A Little Bit Goes a Long Way: Modeling Universal Basic Income for Noncooperative Artificial Agents

Maria Milkowski
Seattle University
Seattle, WA (USA)
mmilkowski@seattleu.edu

Ankur Gupta
Butler University
Indianapolis, IN (USA)
agupta@butler.edu

Nathaniel Kremer-Herman
Seattle University
Seattle, WA (USA)
nkh@seattleu.edu

***Abstract*—With the rise of autonomous, artificially intelligent computer agents, it is broadly accepted that these agents must behave ethically. This requirement is often expressed as a contrast between beneficial, cooperative agent behaviors and detrimental, greedy ones. Many disciplinary approaches have shown that some level of cooperation is beneficial for survival and flourishing. However, it is unrealistic to enforce mass cooperation; some agents (be they real or artificial) will invariably act selfishly.**

Since mass cooperation is fundamentally impractical to enforce, we look for alternative means to achieve similar societal outcomes. We show that a moderate distribution of universal basic income (UBI) is a sufficient surrogate for global cooperation among artificial societies. In particular, we compare five different UBI schemes in the agent-based Sugarscape model and find that artificial societies with the best performing UBI schemes are 5× larger, engage in 6× as much trade, and have agents which live almost 2× as long as in societies without UBI. These results match or exceed those of societies that engage in utilitarian, cooperative behavior patterns.

***Index Terms*—agent-based modeling, economics, machine ethics, computational cultural modeling, cooperative systems.**

I. Introduction

As humanity has become more reliant upon autonomous computer agents to make decisions on our behalf, it has become imperative that these artificially intelligent (AI) agents behave ethically. It is beneficial for AI agents to behave according to well-understood ethical theories (such as utilitarianiasm, Aristotelian virtue ethics, or Confucian ethics) as opposed to letting a machine figure out its own ethical code of conduct or permitting development teams to create *ad hoc* codes of ethics based on their own (incomplete or biased) understanding of ethics. There are numerous benefits to AI agents adhering to ethical theories, but we highlight two: clear explainability of AI ethical reasoning by scholars of ethics and transparent, *a priori* design of ethically-behaving AI agents.

In previous work [1], [2], we discovered the viability of translating a human ethical theory into a computational landscape by interpreting utilitarianism as first introduced by Jeremy Bentham [3] as a decisionmaking algorithm. We wrote this algorithm as the decisionmaking function for autonomous agents in our implementation of the agent-based societal simulation *Sugarscape* [4]. We then tested the degree to which agents (and society) benefit when agents act more selfishly or selflessly. We concluded that the virtual society had the best outcomes across a variety of metrics when agents acted with a perfect balance of selfishness and selflessness.

This egalitarian approach met Bentham's conception of utilitarianism where agents more-or-less democratically made decisions with the agents around them, even if it meant sacrificing themselves or one of their neighbors for the benefit of the rest. This experimentation validated a wealth of literature regarding the nature of greediness and cooperation across a number of disciplines. However, the best societal benefits relied upon an unrealistic expectation: every agent was willing to engage in democratic, egalitarian decisionmaking. In a closed system where all agents' behavior can be pre-programmed by a single authority, cooperation can be enforced. However, like in human societies, autonomous agents are diverse. They come from various designers with their own missions and goals.

With this unrealistic assumption in mind, can we achieve the same (or better) societal results without relying on all AI agents acting cooperatively? To answer this question, we go beyond ethical theory and take inspiration from the social sciences. We experiment with universal basic income (UBI) policies in Sugarscape as a form of digital social welfare to gauge whether some form of digital social safety net in the face of differing agent motives can provide equivalent benefits to mass cooperation. We assume the worst case scenario for our experimentation: all agents act inherently greedily. We measure success by comparing such societies with their utilitarian-minded counterparts; in this sense, UBI societies succeed if their outcomes mirror those of utilitarian societies.

By guaranteeing a UBI safety net to all agents, the computational landscape itself acts as a balance to the greedy, combative agents which traverse it. We show agents in our simulated society are more than 2× as prosperous across most metrics, that society is 5× larger, engages in nearly 6× as much trade, and that agents have a higher degree of happiness in societies with an appropriately-sized UBI than those without. We also note that these digital UBI policies mostly provide superior results to the ideal yet unrealistic case of fully cooperative societies. However, we also caution against overly-large UBI which leads to detrimental outcomes.

II. Brief Overview of Universal Basic Income

We first define UBI as a social policy. UBI has become a popular idea in the social consciousness. American busi-

979-8-3315-3228-4/25/$31.00 ©2025 IEEE

nessman and politician Andrew Yang is often credited for its recent popularization in the United States as a consequence of his book *The War on Normal People* [5]. Yang argues that UBI would increase the quality of life for the general population in a world where job automation is outcompeting human labor and where many careers are increasingly characterized as *bullshit jobs* [6]. The work of Hasdell et al. at the Stanford Basic Income Lab [7] informs our understanding of UBI as having five essential qualities:

1) It is universal, not targeted at a specific population.
2) It is unconditional.
3) It is a cash payment with no stipulations how it is spent.
4) It is paid on an individual basis.
5) It is a recurring payment.

These criteria fit the spirit of the earliest mentions of UBI as we understand the term today. Early American political philosopher Thomas Paine was particularly concerned with the welfare of colonial Americans and was an outspoken proponent of human rights. He laid out the purpose of a UBI policy based on a one-time estate tax in *Agrarian Justice* [8]:

> To create a National Fund, out of which there shall be paid to every person, when arrived at the age of twenty-one years, the sum of fifteen pounds sterling, as a compensation in part, for the loss of his or her natural inheritance, by the introduction of the system of landed property: And also, the sum of ten pounds per annum, during life, to every person now living, of the age of fifty years, and to all others as they shall arrive at that age.

Paine's proto-UBI policy would not become a reality. Though only a proposed system, it meets all of Hasdell's criteria with a gracious interpretation of the *recurring payment* criterion. In Paine's policy, a payment occurs at the age of majority (21) and upon reaching elderly age in Paine's day (50). The recurring interval then becomes annual for the remainder of a person's life after age 50.

More recently in 2018, the Finnish government tested a basic income policy where recipients replaced other forms of social welfare with a direct cash payout [9]. They found virtually no difference in beneficial outcomes, regardless of whether a person received cash or an equivalently-valued service or good. In Finland's case, this would theoretically allow them to dispense with the bureaucratic overhead of continuing their traditional welfare plans.

The US state of Alaska has established a basic income program [10], [11] using the funds from taxation on Alaska's oil and mining industries. This has resulted in lower rates of poverty and deep poverty for Alaskans as well as an increase in family size. The Eastern Cherokee Nation provides unconditional income through their casino dividend [12] with even more impressive results. In the US state of California, a limited scope UBI scheme named SEED found that participants primarily used the income to pay for food (40%), retail needs (24%), utility bills (11%), and auto (9%) expenses. Finally, an OpenResearch UBI study [13], [14], [15] shows how randomly-sampled, low-income individuals in the US states of Texas and Illinois saw improvements in stress, food security, leisure hours, transportation, and finances with only nominal reductions in individual income and labor hours. UBI appeared to give participants the ability to choose more leisure time as well as decide how best to use that time. These real-world instances of basic income would seem to buoy the claimed benefits of UBI policies [5].

Our brief glimpse at basic income policies shows they are typically proposed to address social injustice. We acknowledge these laudable goals, and we seek to extend this conversation by providing a computational perspective to put these best intentions to the test in an AI ethics context. We implement a UBI mechanism adhering to Hasdell's criteria in the societal-scale Sugarscape agent-based simulation. We are not experts in economic policy and implement UBI in Sugarscape using a layperson's understanding of the policy. Through our Sugarscape model, we observe the following societal outcomes of various UBI schemes:

1) Societies with a small amount of UBI have increased societal prosperity across a variety of metrics.
2) Individuals in societies with a small amount of UBI are happy across a variety of definitions of happiness.
3) Societies with too large an amount of UBI have many detrimental outcomes at the societal level.

III. Related Work

Simulation is a long-standing method of investigating the validity and usefulness of various social policies. Their use often comes in the form of mathematical models and microsimulations [16], [17], [18]. They have demonstrated the difficult yet crucial task of forecasting future policy changes or radical shifts in societal conditions without real-world data from these prospective futures [19].

There are relatively few agent-based models (ABMs) for social policies when compared to the mathematical modeling and microsimulation approaches though they have been used equally as long for social simulation. Gilbert and Doran's exploration of ABMs [20] in the 1990s heralds the firm roots of agent-based computer simulation in the social sciences. ABMs expose the micro-level agent behaviors and interactions that provide richness and dynamicity to complex social simulations [21] which the microsimulation approach cannot effectively capture. ABMs allow individual agents to express preferences and hold normative values [22], [23].

Recent agent-based models have demonstrated their usefulness exploring social policies. For example, agent-based modeling has been applied to police funding schemes to determine their potential impacts on crime rates [24] more rapidly than policymakers could deliver through real-world directives. Automation as a consequence of innovation has a significant effect upon wages and employment. Neves et al. [25] use an ABM to explore this, underscoring the context of Yang's [5] discussion of basic income.

The Sugarscape ABM we use for our simulation framework was first introduced in *Growing Artificial Societies* [26] as a simulation for computational social sciences. Sugarscape has been used to study many social phenomena given its societal-level insights about emergent behavior. Perhaps most relevant to this work is an investigation into the benefits of social security and pensions using Sugarscape [27]. Other works have used Sugarscape to examine tax structures [28], wealth adjustment and disparities [29], and the proliferation of technology as it relates to societal formation [30]. While Sugarscape was first introduced in the 1990s, its contemporary use demonstrates its continued usefulness as a viable model for computational social science.

Sugarscape has also been applied to ethics. One attempt at translating human ethics into Sugarscape agents was based on Aristotelian virtue ethics [31]. The virtue of temperance was represented as a means of moderating agent consumption. Our previous work [1], [2] translates utilitarianism, quite a different ethical theory, into a decisionmaking algorithm. Both these bodies of work contribute to the field of machine ethics which is concerned with how we create ethically-behaving AI agents. There is still significant work to do regarding using ABMs to represent ethical issues [32], and ABMs are only nascently being used to model explicitly ethical AI agents.

IV. Sugarscape

Sugarscape models societal-level emergent behaviors from individual agent behaviors. The simulation environment is a two-dimensional $n \times m$ toroidal grid. Two resources exist across the environment: *sugar* and *spice*. Grid cells have a (potentially zero) initial allocation of both resources, and cells regenerate up to this initial allocation over time. Based on a user-provided configuration, the environment may be particularly hostile or hospitable to life.

An initial population of agents is randomly spread around the environment. Agents are born inherently greedy and aim to live as long as possible. Each agent has a *metabolism* for sugar and spice, and they consume their metabolism's worth of both resources per simulation timestep. They may also move to a new cell according to their *vision* to gather more sugar and spice in the environment. Agents may perform a host of other behaviors (if enabled) such as trading, lending, reproducing, making friends, murdering neighbors, passing along diseases, and exerting cultural pressure upon others.

Sugarscape is a fantastic sandbox for investigating agent-based behavior as it demonstrates rich social dynamics arising from simple agent rules. While it is coarse-grained compared to other societal-scale simulations, its strong metaphors for the real world are attainable through discrete numerical values that are readily computable [26]. This introductory overview provides the necessary context to understand our experimentation. A more complete discussion of this Sugarscape implementation can be found in our previous work [4].

The top left of Figure 1 demonstrates an initial random spread of agents (red), sugar (yellow), spice (brown), and cells with both resources (tan). The greater the saturation of color at a cell, the more resources are present at that cell. The top right shows the society 75 timesteps into the simulation: an initial population dip has occurred due to competition, disease, and natural selection. The bottom left shows societal rebound 125 timesteps into the simulation. Finally, the bottom right shows a prosperous, successful society at timestep 200.

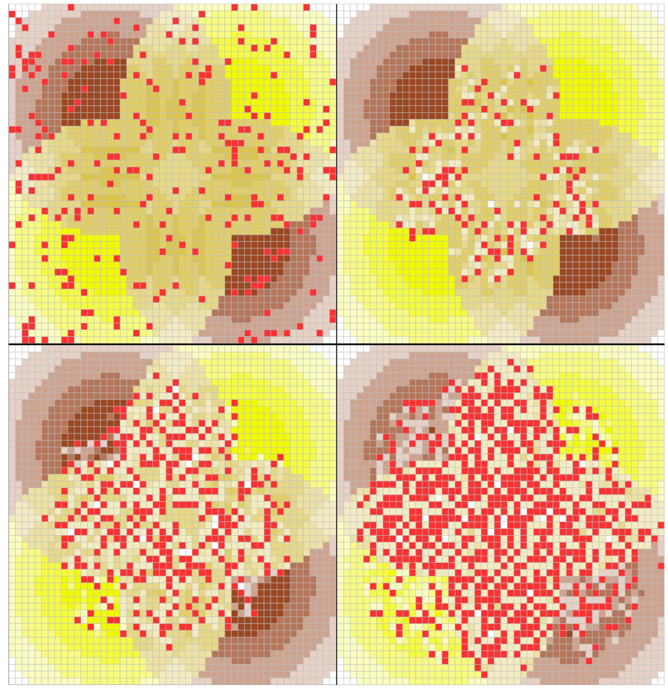

Fig. 1. A Sugarscape Society from Beginning to End

A. Universal Basic Income in Sugarscape

Each simulation timestep, every agent receives a UBI payment on their turn. In keeping with Hasdell's criteria [7], our implementation of UBI in Sugarscape is universal, unconditional, individual, and recurring. While the free sugar and spice are not a cash payment *per se*, the agent receives raw resources which are an unrestricted cash-equivalent in Sugarscape: it can be consumed, traded, lent out, or used to pay outstanding loans at the agent's prerogative. Sugarscape can be configured to provide UBI with variable amounts of sugar and spice, paid at specified intervals of timesteps, and allows for unequal payouts across agents. However, we do not enable these more complex distribution methods or patterns. We focus on configurations that match Hasdell's criteria giving each agent the same payment every timestep.

B. Agent Happiness in Sugarscape

A primary claim for Yang is that UBI leads to greater happiness [5]. We model happiness along different axes in Sugarscape testing to what degree UBI contributes to that happiness. We represent each component of agent happiness distinctly as well as provide a composite happiness score. An agent experiences (un)happiness along conflict, family, health, social, and wealth axes. Each of these values is assessed per timestep and ranges from -1 to 1.

Conflict happiness is decreased by 1 by killing another agent (representing remorse). Some agents can be configured to gain happiness in combat, but this is not enabled in our experimentation. **Family** happiness is increased according to the size of an agent's living family (mates and children) and mating. They lose happiness according to their deceased family members. **Health** happiness is 1 if an agent is healthy but -1 is the agent is sick. **Social** happiness is increased by having more friends in their social network but lose happiness when they have few or no friends. **Wealth** happiness is increased with having wealth (combined sugar and spice) higher than the global average. They lose happiness when their wealth is lower than average. For family and wealth happiness, the gain or loss is based on the Gauss error function (commonly called *erf*). For social happiness, the value is computed by a stepwise function similar in nature to *erf* limited by the agent's configured maximum number of friends.

To compute the *composite happiness* of an agent, we sum the individual happiness measures. Therefore, composite happiness ranges from -5 to 5. If the hypothesized benefits of UBI policies are realized in Sugarscape, we would expect agents to have a happiness closer to 5. Note that, at best, an agent that does not gain happiness through combat will experience a total happiness of 4.

V. Evaluation

We evaluate the feasibility of various UBI policies as substitutes for enforcing global cooperation in Sugarscape. We take a worst-case approach to evaluation by allowing agents to act uninhibitedly greedily. To determine the impact of any particular policy, we look at societal-level outcomes. While it may be tempting to focus on micro-level details of particular agents, in reality these social policies are implemented to provide benefits at the macro-level. As such, we only consider metrics which tell us about societal-level changes.

Since Sugarscape has two resources (sugar and spice), we configured the simulation to provide UBI for both. Our dataset includes combinations of sugar and spice basic income across the following values: 0, 1, 2, 3, 4, and 10. A policy of zero sugar and zero spice is the control case (no UBI). In all experiments, agents have distinct metabolisms for both sugar and spice which could range between 1 and 4 (hence the range in our experimental UBI values). We added the final value, 10, to our experimental range to see what would happen if the UBI far exceeded the metabolic needs of any agent.

We ran all combinations of UBI schemes across 200 random seeds (randomized starting states for the simulation). Since Sugarscape is deterministic, different outcomes between UBI policies for the same seed are truly indicative of direct effects of the policy. Each simulation run lasted 5,000 timesteps (or until societal extinction) and began with an initial population of 250 agents. Agents were configured to have a maximum age ranging from 60 to 100 timesteps, so the simulation ran for 50 generations. In our experiments, nearly all simulation options are enabled most importantly including: aging, combat, cultural tribes, disease, lending, reproduction, and trade. More details about these features are provided in previous work [4].

A. Sugar and Spice Basic Income Schemes

Our results include UBI schemes with a matched amount of sugar and spice and are expressed in the form of a tuple $(sugar, spice)$ and ranges from $(0, 0)$, representing no UBI, to $(10, 10)$, representing 10 sugar and 10 spice income per timestep. In Sugarscape, both resources are combined to calculate an agent's wealth. This single wealth score is used for important actions such as determining their strength in combat and for determining which cells best meet their needs during movement. Since this combined wealth is used throughout the simulation, it makes sense to consider a UBI policy for Sugarscape in which both sugar and spice are equally valued.

For completeness, we also collected results for mismatched UBI schemes where, without loss of generality, less spice is provided than sugar such as $(2, 1)$ or $(3, 2)$ and only sugar UBI schemes such as $(1, 0)$ or $(10, 0)$. The results for these mismatched schemes cluster with the results for their matched counterparts. For example, the $(3, 0)$ and $(3, 1)$ schemes cluster with the $(3, 3)$ scheme. Hence, they are excluded in the discussion of our results. However, the simulation runs are contained within our dataset and are available for replication.

VI. Results

Across a variety of commonsense metrics, a modest UBI provides strong benefits which perform as well or better than globally cooperative, utilitarian behavior. There are also detrimental effects of providing increasingly large amounts of basic income which lead to poor societal outcomes. In the case of UBI, a little truly goes a long way in Sugarscape.

TABLE I
Societal Survival

Behavior	UBI Scheme	Extinct	Worse	Better
Cooperative	$(0, 0)$	2	0	198
Greedy	$(0, 0)$	69	0	131
Greedy	$(1, 1)$	2	1	197
Greedy	$(2, 2)$	10	2	188
Greedy	$(3, 3)$	56	5	139
Greedy	$(4, 4)$	78	9	113
Greedy	$(10, 10)$	83	8	109

A. Societal Survival

The most direct gauge of societal success is survival. If a policy leads to collapse, there is no use looking at any further metrics of success. Table I shows the survival rate of societies across various UBI schemes compared to our gold standard of mass cooperation. We compare the ending population with the beginning to sort each run into one of three buckets: extinct (zero population at the end), worse (fewer than or equal to the starting agents), and better (strictly greater agents at the end). The greedy $(0, 0)$ case presents the default behavior of the simulation while the cooperative $(0, 0)$ case uses the utilitarian decisionmaking algorithm presented in previous work [1].

A poor initial configuration can lead to tough starting conditions with rampant competition. There is a high degree of volatility due to the greedy nature of agents in the $(0, 0)$ control case. Roughly a third of the societies led to extinction.

This is contrasted with the near-perfect flourishing of the ethical, cooperative approach. By caring for the needs of other agents, each agent in the simulation eschews claiming a resource-rich cell if it would help another nearby agent more. The two extinct societies in spite of mass cooperation are indicative of especially harsh initial conditions.

Providing even $(1, 1)$ UBI boosted survival dramatically from the control and is practically equivalent to mass cooperation. With a modicum of UBI, the greedy society is as well-off as the fully cooperative society. The two cases of extinction and single case of society being worse off at the end are due to harsh initial conditions. One would think increasing UBI would correct the few remaining poor outcomes.

However, providing a greater matched UBI than $(1, 1)$ led to worse societal survival. The $(2, 2)$ scheme provides a respectable increase from the control, but it falls below the $(1, 1)$ scheme and leads to slightly more extinction. Taken alone, this could indicate the $(2, 2)$ scheme is worse than $(1, 1)$. However, we later share metrics that reveal some redeeming qualities of the $(2, 2)$ scheme despite its lower societal survivability.

With even greater UBI, the benefits erode. The $(3, 3)$ scheme performs only slightly better than the $(0, 0)$ control, but the $(4, 4)$ and $(10, 10)$ schemes are significantly worse. The $(4, 4)$ scheme represents all agents having their metabolic needs met per timestep. We will later show that high UBI leads to high levels of endemic disease in Sugarscape society.

B. Population

Population size over time is a straightforward metric to measure. According to UBI proponents, such policies encourage greater reproduction in society [5]. In Sugarscape, the initial population of agents are randomly placed on the environment with a starting amount of sugar and spice. From then on, the only new agents to populate the grid are born from sexual reproduction. This reproduction requires both parents to contribute a portion of their resources (which must be at least as much as their starting sugar and spice). This birth endowment given to the child represents critical resources needed for child rearing.

Figure 2 shows the mean population across all 200 seeds for all UBI schemes and the cooperative approach across 5,000 simulation timesteps. The environment is a 50 × 50 grid, and no more than one agent can occupy a cell. So, the maximum population is 2,500 agents. In case of societal extinction, a zero value was propagated for the remainder of the timesteps out to 5,000. While this depresses the values for cases with many extinctions, the relative order of performance does not change. We continue this zero propagation in our results.

The $(0, 0)$ control case is only marginally successful at fostering large societies. There is an initial *murderous period* at the beginning of each simulation run where agents scramble for the best resources evidenced by the initial population dip

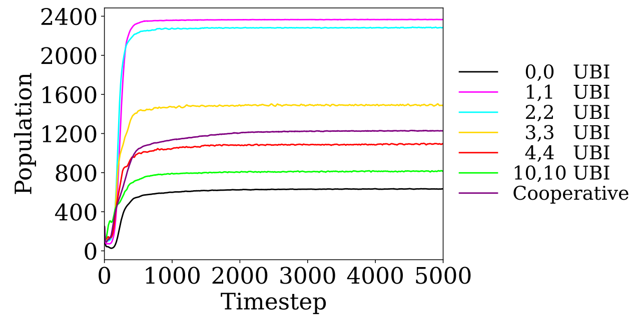

Fig. 2. Mean Societal Population

across results. Should a society survive this dip, they are likely to make it to the end of the simulation. For the $(0, 0)$ scheme, this dip is the most severe and depresses the ability for society to bounce back in the long term.

The cooperative approach results in nearly 3× as large of societies compared to the $(0, 0)$ control. The murderous period dip is not as severe. By caring for the wellbeing of their neighbors, agents ensure that their neighborhoods more optimally allocate resources. Additionally, poorer agents are given the opportunity to claim higher-value cells which provides more resources for their own survival and for reproduction.

The $(1, 1)$ scheme has a shallow dip during the murderous period and is the best performer overall. The small amount of UBI granted is enough to sharply decrease the barrier for starting (or expanding) a family. However, this small UBI will only fully satiate the metabolic needs of the most fit agents (who have 1 sugar and 1 spice metabolism per timestep). As a result, the vast majority of the population is reliant upon the environment's resources for survival.

Other schemes perform better than the $(0, 0)$ control but worse than the $(1, 1)$ policy. The $(2, 2)$ scheme is a close second to $(1, 1)$ but with worse societal survival. The $(3, 3)$ policy performs far worse than $(2, 2)$. This trend continues with the $(4, 4)$ and $(10, 10)$ schemes, even though these schemes fully supply the metabolic needs of all agents. These two schemes also perform worse than the cooperative case.

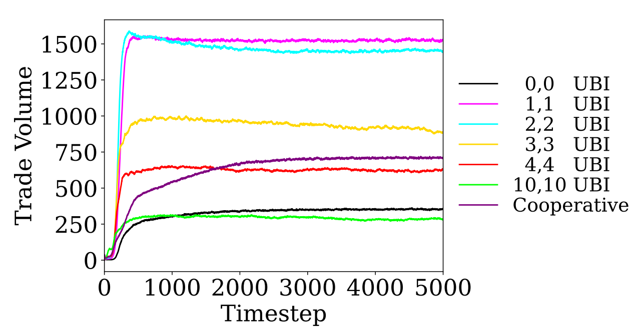

Fig. 3. Mean Societal Trade Volume

C. Trade

One of the purported benefits of a basic income policy is increased spending. The extra cash for citizens incentivizes greater spending on necessities and luxuries. While Sugarscape has no truly analogous spending mechanism (though reproduction costs can serve a partial role), agents can trade one resource for another. When an agent moves to a new cell, they may conduct trades with one or more of their new neighbors. The trade volume shown in Figure 3 is a mean of the tally of trade transactions across all seeds across all timesteps.

We validate this hypothesized benefit of UBI as nearly all schemes lead to significant increases in trade. The $(1, 1)$ and $(2, 2)$ policies show clear benefits compared to the other policies. By providing UBI which will not satisfy the metabolic needs of most agents, they must gather and trade resources to survive. They also greatly outperform the cooperative approach which is already more than $2\times$ as good as the $(0, 0)$ control. However, the larger UBI policies demonstrate the detrimental effect of providing too many resources: the $(10, 10)$ scheme does such a poor job motivating agents to trade that it performs as badly as the control.

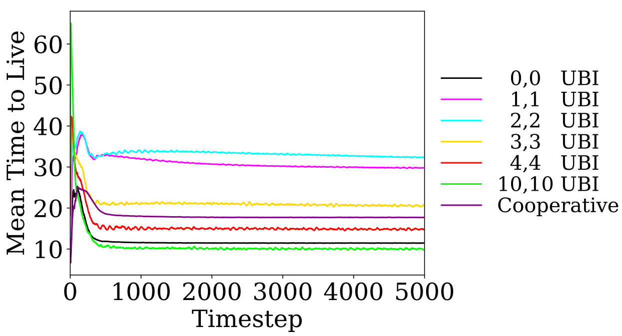

Fig. 4. Mean Agent Time to Live

D. Life

Time to live (TTL) is a measure of how many timesteps an agent can survive with their metabolic needs given their current sugar and spice stockpile without acquiring any additional resources. We also include an agent's UBI, since this income is guaranteed. The maximum number of timesteps for an agent's TTL is their pre-configured maximum age. A higher societal mean TTL is an unqualified good, especially since the maximum age for agents is a fairly constrained range of values.

Figure 4 shows the mean TTL of agents across all seeds and all timesteps. The $(0, 0)$ control results in the mean agent holding resources to sustain them roughly 10 timesteps. By simply cooperating with other agents, however, mean TTL nearly doubles. One would expect with basic income which actively provides extra resources, agents' TTLs must surely increase. However, we see a counter-intuitive pattern of diminishing returns, much like the population results.

The smaller UBI schemes vastly outperform the control. The $(2, 2)$ scheme edges out the $(1, 1)$ scheme in mean TTL as agents are able to bank more resources for leaner times. This observation fails, however, as the UBI continues to increase. Once basic income becomes large enough, agents have bleaker horizons to their lives. The $(10, 10)$ case, which one might expect to have the highest performance, performs worse than the control. We show later this is due to endemic disease.

Agents are provided enough resources to cover many of their necessities in the $(1, 1)$ and $(2, 2)$ schemes. As a result, mean TTL is roughly half an agent's expected lifetime. This incredible safety net practically guarantees an agent survives harsh periods where they are unable to gather many resources over a long stretch of timesteps.

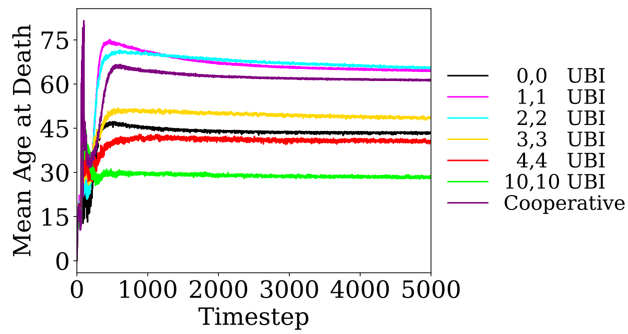

Fig. 5. Mean Agent Age at Death

E. Death

When considering mean age at death, the full tragedy of the degenerate effects of too much UBI become apparent. Figure 5 shows the mean age at death for agents across all seeds for all timesteps. The $(10, 10)$ policy leads to a disappointingly short life expectancy of 30 timesteps. While agents are not dying left and right, they are dying astonishingly young. With a life expectancy of 42 timesteps, agents in the control case are likewise short-lived. The $(4, 4)$ scheme performs worse than the control but is mollified in its harmful effects since payments are much smaller than the $(10, 10)$ scheme.

The remainder of the cases improve upon the $(0, 0)$ control. The $(1, 1)$ and $(2, 2)$ cases again stand out as strong performers, narrowly edging out the cooperative case. They provide enough to maximize life expectancy without leading to detrimental outcomes. By the end of the simulation, these two policies perform practically the same with a slight edge for the $(2, 2)$ policy demonstrating it is closer to an ideal provisioning of resources to maximize the life expectancy of any arbitrary agent in society.

F. Disease

A naïve analysis might lead one to consider large basic income policies would lead to significant societal prosperity. However, such policies lead to diminishing returns or detrimental outcomes which worsen prosperity. In Sugarscape, it is endemic disease that often becomes the degenerating effect upon societal success. The results for population, trade, mean TTL, and mean age at death are tarnished by unrestrained disease propagation in societies with larger UBI schemes.

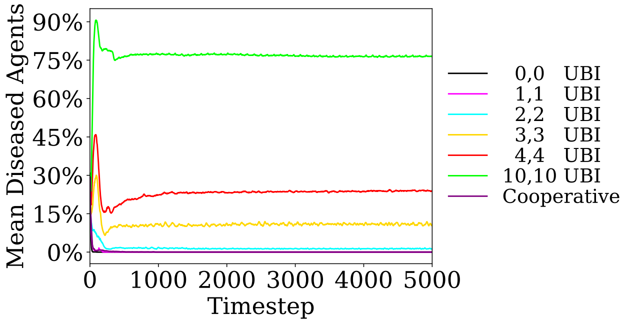

Fig. 6. Mean Diseased Agents

Each society is configured to begin with 50 randomly generated diseases. These affect an agent's metabolism, vision, reproductive capability, aggression, and more. An agent can contract many diseases, and agents slowly (attempt to) recover from them over time as their immune system adapts. More details on the disease and immune system mechanisms are provided in *Growing Artificial Societies* [26].

After an agent moves, they randomly communicate a disease to each of their neighbors. In tightly clustered societies, this leads to most agents having multiple diseases which lead to higher sugar and spice metabolisms. With a high enough UBI, the initial few increases to an agent's metabolic rates are offset by the free resources per timestep. This keeps agents with incredibly pernicious diseases from dying (instead leading poor-quality lives) long enough to become superspreaders of these diseases. Eventually, these illnesses become endemic. The reduced population creates few offspring who soon catch too many diseases to effectively recover from any (as each acts upon the immune system mechanism independently).

Figure 6 shows the percentage of the population with at least one disease across all seeds and all timesteps. While there are other contributing factors for why some policies perform better or worse than others, the context of disease provides the rest of the story. The higher the UBI, the more often disease becomes endemic in Sugarscape societies.

Disease is practically eradicated in the $(1, 1)$ UBI and cooperative cases while also leading to strong societal benefits. In the $(0, 0)$ control, diseased agents are violently culled in the murderous period. In the $(1, 1)$ scheme, these unfortunate agents are provided compassionate, palliative care through their UBI without being forced to continue a tortured existence (and spreading their illness across society).

The $(2, 2)$ UBI scheme performs quite similarly to $(1, 1)$ UBI, the cooperative approach, and the $(0, 0)$ control. Given the relative benefits the $(2, 2)$ scheme has over $(1, 1)$ UBI in mean TTL and mean age at death, the small (but not insignificant) increase in endemic disease may be worth the gains. For the larger policies, the outcomes degenerate rapidly because endemic disease rates increase sharply. The $(10, 10)$ UBI scheme stands out as a horrific result: about 75% of society carries at least one illness. Given the abysmal mean TTL results, it is likely these agents carry multiple diseases.

G. *Happiness*

UBI policies are purported to increase the happiness of the citizenry. We model agent happiness along five axes: conflict, family, health, social, and wealth happiness. The sum of these five scores form a composite happiness score for each agent. We expect basic income policies to lead to higher happiness in each category. Table II shows a condensed view of the various kinds of happiness for each case under study. We provide the mean values at timestep 5,000 for each type of happiness.

Wealth happiness is negative across the board since all societies experience wealth disparity, however the amount of unhappiness is surprisingly greatest in the $(1, 1)$ and $(2, 2)$ schemes. While not included in our discussion, this is due to these two schemes having a higher level of wealth inequality. Since no agent was configured to derive happiness from combat, conflict happiness settles to zero at the end of the simulation for all cases. Family happiness overall provides a small bonus as not all agents have families, and this bonus becomes noticeably smaller in the volatile control case and the high UBI cases. Health happiness is also high in societies with high success rates, especially since societies that survive most often eradicate disease. Likewise, societies with high survival rates also provide a stable setting for agents to make friends which leads to higher social happiness.

In total, the $(1, 1)$ and $(2, 2)$ UBI schemes perform exceedingly well across metrics and also result in the highest mean composite happiness scores among UBI policies. These success stories suggest that a modest UBI is a more ideal allocation of resources. Providing more resources brings societies over the edge, resulting in progressively lower happiness even when compared to the control case.

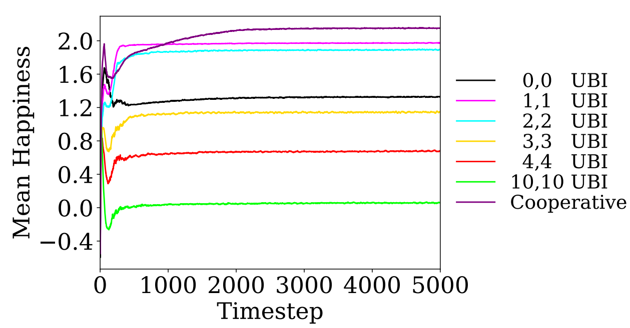

Fig. 7. Mean Agent Composite Happiness

Figure 7 shows the mean composite happiness across all seeds and all timesteps. The murderous period results in a significant dip in overall happiness, however societies quickly adapt and settle into stratified scores. The cooperative approach results in the highest overall happiness closely followed by the two best performing UBI schemes. Every other UBI scheme results in less happiness than the control with the $(10, 10)$ case leading to nearly negative mean happiness.

For the $(4, 4)$ and $(10, 10)$ policies, societal happiness post-dip does not recover to nor go beyond the pre-dip high. Like most metrics, the larger UBI policies have lower overall

TABLE II
MEAN HAPPINESS METRICS AT 5,000 TIMESTEPS

Behavior	**UBI Scheme**	**Conflict**	**Family**	**Health**	**Social**	**Wealth**	**Composite**
Cooperative	$(0, 0)$	0.00	0.31	0.99	0.95	-0.10	2.15
Greedy	$(0, 0)$	0.00	0.21	0.66	0.62	-0.16	1.33
Greedy	$(1, 1)$	0.00	0.27	0.99	0.97	-0.26	1.97
Greedy	$(2, 2)$	0.00	0.27	0.89	0.93	-0.20	1.89
Greedy	$(3, 3)$	0.00	0.20	0.42	0.69	-0.17	1.15
Greedy	$(4, 4)$	0.00	0.16	0.09	0.58	-0.14	0.68
Greedy	$(10, 10)$	0.00	0.10	-0.41	0.51	-0.14	0.06

happiness compared to the control case because they have provided more resources than what is beneficial for society. The smaller UBI schemes come much closer to this ideal basic income and result in a far higher mean composite happiness.

The small UBI policies therefore result in far greater happiness compared to the control and larger UBI schemes. Across the metrics considered, these smaller payments approach an optimal resource allocation which maximizes the purported benefits of UBI (greater population, more economic activity, increased life expectancy, and greater happiness). Basic income policies that provide resources beyond this optimal allocation see rapidly diminishing returns.

Even without free resources, the cooperative agents experience greater happiness than all other schemes. Table II provides context for why the cooperative approach leads to greater happiness. While a society entirely populated by considerate agents provides many benefits, it has an (unfortunately) vanishingly small possibility of occurring in real life.

In the metrics shown, the cooperative societies vastly outperform the control case. However, providing a UBI policy provides at least as strong an outcome as cooperative societies. Only in happiness, disease, and overall societal survival are there minor caveats. In this sense, UBI can serve as a surrogate of sorts. Since mandating ethical social behavior for all agents is untenable, providing UBI to achieve similar societal success is the feasible alternative.

VII. DISCUSSION

Our results validate the benefits of UBI policies purported by proponents [5]. Sugarscape agents receiving a small UBI have far more children, live much longer, and engage in greater amounts of trade. Since the simulations ran for roughly 50 generations, these benefits are demonstrated on a timescale impractical to replicate with real-world UBI experiments. Our results also reinforce that Sugarscape, with its rich agent interactions and emergent societal outcomes, is a useful agent-based model for simulating social policies generally.

A. Broader Applicability

Our primary concern is whether the established benefits of UBI can be translated to artificial agents. We demonstrate this is the case and directly build upon our previous work [1] establishing the benefits of mutual cooperation for AI agents. Since autonomous agents may come from a variety of manufacturers, serve various missions, and may be programmed with a diverse assortment of ethical principles, we cannot enforce strong guarantees of globally cooperative behavior.

So, how might one apply the demonstrated benefits of UBI in an arbitrary autonomous agent context? In Sugarscape, we compensated for the lack of ethical behavior from the agents by providing a safety net from the computational environment. In like manner, one can create a safety net in any arbitrary system filled with autonomous agents so long as there is some known, interactive mechanism similar to the Sugarscape environment which agents engage periodically. Further, as demonstrated in our results, there is an optimal provisioning of resources (in our experimentation, just a little bit extra) which helps the system. Beyond that optimal provisioning, stagnation or destabilization may occur.

Since UBI provides free resources to agents, such a mechanism is only applicable in applications where some universal provisioning of extra resources is feasible. Smart environments upon which autonomous agents act are ideal candidates for the UBI approach. For example, a smart traffic system is an environment (with smart traffic lights and a plethora of sensors) upon which a diverse range of self-driving vehicles operate. The smart traffic system might provide a token-based UBI such that cars programmed to drive more defensively can exchange tokens in the system to avoid systematic exploitation by cars programmed to drive more unscrupulously.

Monolithic AI agent systems, where there is a single agent with no interaction with others, cannot benefit from this UBI approach. A web browser chatbot powered by a large language model, for instance, does not directly interact with any other chatbots. We cannot rely on a computational environment or other autonomous agents to keep the chatbot in check via resource exchange. Instead, we must insist that these chatbots intrinsically adhere to well-understood ethical principles as argued in our initial work [1], [2].

B. Limitations

Sugarscape takes an abstract view of society. This coarse-grained approach provides key benefits such as inherent computability and a greater focus on emergent societal-level behavior. Our implementation of UBI in Sugarscape is likewise abstract which provides a few specific limitations to our study. Each of these limitations provide opportunity for future work to further increase the usefulness of Sugarscape as an experimental framework for social policies.

1) Limited Conceptualization of Agent Happiness: Our conceptualization of agent happiness is defined along five axes. These different kinds of happiness stem from the actions an agent may (or may not) take during their turn each timestep (moving, making friends, trading, lending, reproducing, exerting cultural influence, killing, and spreading disease). These actions are but a subset of what people can do.

Our calculation of the happiness values are straightforward. They are intended to be quickly computed and provide a fair look into the state of any given agent. While a great start, there is room for improvement. Some happiness metrics are quite simple, such as conflict and health happiness. We also assumed (for sake of computability and comparison) that (un)happiness of any given type could be bound between a range of -1 and 1 representing profoundly unhappy and profoundly happy states. This assumption may not be realistic enough if one were to use agent happiness as the sole evidence that UBI policies will lead to increased prosperity.

2) Limited Conceptualization of Agent Motivation: We also have a limited conceptualization of an agent's motivations. Sugarscape was originally designed with agents having one motivation: live for as long as possible by collecting resources in their environment [26]. Our reimplementation of Sugarscape [4] expands this notion by introducing other motivations for agent actions (such as mutual cooperation).

In all cases these Sugarscape agents behave as *Homo economicus* [33]. We assume agents always act rationally; this rationality is driven by resource gathering as in the original implementation. It remains future work to expand agent motivation to include non-material reasons to act (such as tightening familial bonds) which are abound in human behavior. This phenomenon is a common critique of economic modeling to which Sugarscape also falls prey. However, the insights we provide are of *limited* value compared to a true-to-life simulation (which is beyond computational capability), rather than it being of *no* value.

VIII. Reproducibility of Experimentation

Because the Sugarscape implementation is deterministic, our results are exactly reproducible.[1] Software requirements are provided in the `README` of the software repository. The results provided in this work were created using the `v2024.3` release. Using another version will produce similar, but not exactly the same, results. We provide the full set of JavaScript Object Notation (JSON) configuration files necessary to reproduce our findings.[2] The configurations are located in the `little-bit-ubi.zip` archive.

The simulation takes as input a single configuration file. Once extracted, each of the JSON configuration files can be re-run one-at-a-time. To run a single iteration of Sugarscape with a configuration file, run the following (substituting the system's proper Python 3 alias):

```
> python sugarscape.py -c <CONFIG>
```

[1]Software at: https://github.com/digital-terraria-lab/sugarscape.

[2]Configuration files at: https://github.com/digital-terraria-lab/datasets.

To re-run the full dataset, the top-level software repository configuration file `config.json` must be modified. The `numParallelSimJobs` property specifies how many simulation instances are allowed to run in parallel locally. This should be scaled to the number of CPU cores or hardware threads available on the machine leaving enough resources behind to effectively run the operating system. Sugarscape is memory-efficient, so the bottleneck for data collection is available cores. The dataset configuration files should be placed in the repository's `data` directory before execution.

Before performing data collection, run `make setup` to perform a low-effort check for a local installation of Python 3. This will set values in the `Makefile` and `config.json` accordingly. Running `make data` performs the data collection, and `make plots` will generate graphs provided `matplotlib` is available. We note tinkering with the plotting script in the `plots` directory is necessary to produce exact graphs to those presented in this work due to formatting differences. To generate the dataset and produce graphs, run:

```
> make setup
> make data
> make plots
```

Since local data collection may be slow, the repository can be easily extended for distributed execution. This work used the Open Science Grid [34] and the HTCondor batch system [35] to perform data collection. We used the Makeflow scientific workflow management system [36] to structure the data collection passed to HTCondor. Regardless of batch system used, a script can be written to handle job creation, submission, and management. We encourage this approach when reproducing our results.

IX. Conclusions and Future Work

We investigated the viability of using a universal basic income mechanism as a means of attaining strong societal success across a variety of commonsense metrics in the agent-based social simulation Sugarscape. These benefits serve as a worthy proxy of the results in our previous work [1] which relied on mass cooperation of all agents adhering to a utilitarian decisionmaking algorithm. However, our results make no such assumption; instead, all agents acted greedily.

Our results validate existing literature on basic income policies. A little UBI goes a long way, but providing too large of basic income yields diminishing returns. There is a tipping point where society begins to experience detrimental outcomes rather than simply fewer gains. In Sugarscape, this was exhibited through rampant, endemic disease which led to societies with less prosperity than the control case.

In future work, we would like to address agent motivation. Agents currently only act according to economic priorities (i.e. acquiring sugar and spice). However, there is more to life than acquiring wealth. We foresee powerful changes to the simulation's emergent behaviors once agents begin selecting actions based not only on economic gain but also improving the prosperity of their family, friends, and neighbors. For example, quarantining when sick is not economically beneficial for the acting agent, but this is commonly recognized as socially-conscientious behavior which ought to be included in our simulated societies.

Further, we intend to explore various other ethical principles translated into decisionmaking algorithms. While the utilitarian, cooperative approach presented here can be approximated (or improved upon by) UBI, other frameworks may lead to greater prosperity for agent societies. One related work has provided an initial look at Aristotelian virtue ethics in Sugarscape [31], however there is a diverse range of ethical perspectives stretching across the globe and across time worthy of exploration in a computational landscape.

X. Acknowledgments

We acknowledge Hadiya Chishti, Colin Hanrahan, Willem Hueffed, Anna Muller, Joshua Palicka, Mariana Shuman, Abi Sipes, and Lucas Vorkoper for their contributions to our software repository. We also acknowledge Eric Severson for his guidance and critiques of our machine ethics implementation. This work was supported in part by the Seattle University Undergraduate Summer Research Award and the Holcolm Award Research Fund at Butler University. This work was facilitated in part by using services provided by the Open Science Grid Consortium, which is supported by National Science Foundation awards #2030508 and #1836650.

References

[1] N. Kremer-Herman and A. Gupta, "The need for greed in artificial decisionmakers," in *International Symposium on Technology and Society*. IEEE, 2024.

[2] N. Kremer-Herman, A. Gupta, and E. R. Severson, "Blueprints for machine ethics: A digital terrarium for socio-ethical artificial agent decisionmaking," *IEEE Access*, 2024.

[3] J. Bentham, "An introduction to the principles of morals and legislation," *London: Athlone*, 1789.

[4] N. Kremer-Herman and A. Gupta, "Replacing sugarscape: A comprehensive, expansive, and transparent reimplementation," in *International Conference on Simulation Tools and Techniques*. Springer, 2023, pp. 79–92.

[5] A. Yang, *The war on normal people: The truth about America's disappearing jobs and why universal basic income is our future*. Hachette UK, 2018.

[6] D. Graeber, *Bullshit Jobs: A Theory*. Simon & Schuster, 2018.

[7] R. Hasdell, "What we know about universal basic income: a cross-synthesis of reviews," *Stanford Basic Income Lab*, 2020.

[8] T. Paine, *Agrarian Justice*. R. Folwell, 1797.

[9] O. Kangas, S. Jauhiainen, M. Simanainen, and M. Ylikännö, "The basic income experiment 2017–2018 in finland: Preliminary results," *Reports And Memorandums Of The Ministry Of Social Affairs And Health*, 2019.

[10] M. Berman, "Resource rents, universal basic income, and poverty among alaska's indigenous peoples," *World Development*, vol. 106, pp. 161–172, 2018.

[11] D. Jones and I. Marinescu, "The labor market impacts of universal and permanent cash transfers: Evidence from the alaska permanent fund," National Bureau of Economic Research, Working Paper 24312, February 2018.

[12] I. Marinescu, "No strings attached: The behavioral effects of u.s. unconditional cash transfer programs," National Bureau of Economic Research, Working Paper 24337, February 2018.

[13] S. Miller, E. Rhodes, A. W. Bartik, D. E. Broockman, P. K. Krause, and E. Vivalt, "Does income affect health? evidence from a randomized controlled trial of a guaranteed income," National Bureau of Economic Research, Working Paper 32711, July 2024.

[14] E. Vivalt, E. Rhodes, A. W. Bartik, D. E. Broockman, and S. Miller, "The employment effects of a guaranteed income: Experimental evidence from two u.s. states," National Bureau of Economic Research, Working Paper 32719, July 2024.

[15] A. W. Bartik, E. Rhodes, D. E. Broockman, P. K. Krause, S. Miller, and E. Vivalt, "The impact of unconditional cash transfers on consumption and household balance sheets: Experimental evidence from two us states," National Bureau of Economic Research, Working Paper 32784, August 2024.

[16] A. J. Auerbach and L. J. Kotlikoff, "Evaluating fiscal policy with a dynamic simulation model," *The American Economic Review*, vol. 77, no. 2, pp. 49–55, 1987.

[17] R. R. McDaniel, R. S. Sullivan, and J. R. Wilson, "A simulation model for welfare policy analysis," *Socio-Economic Planning Sciences*, vol. 22, no. 4, pp. 157–165, 1988.

[18] A. V. D. Luduvice, "The macroeconomic effects of universal basic income programs," *Journal of Monetary Economics*, p. 103615, 2024.

[19] A. A. Zagonel, J. Rohrbaugh, G. P. Richardson, and D. F. Andersen, "Using simulation models to address "what if" questions about welfare reform," *Journal of Policy analysis and Management*, vol. 23, no. 4, pp. 890–901, 2004.

[20] G. Gilbert and J. Doran, *Simulating Societies: the computer simulation of social phenomena*. UCL Press, 1994.

[21] F. Squazzoni, W. Jager, and B. Edmonds, "Social simulation in the social sciences: A brief overview," *Social Science Computer Review*, vol. 32, no. 3, pp. 279–294, 2014.

[22] R. Mercuur, V. Dignum, and C. M. Jonker, "The value of values and norms in social simulation," *Jasss*, vol. 22, no. 1, p. 9, 2019.

[23] A. Melnyk, B. Edmonds, A. Ghorbani, and I. van de Poel, "Modelling values in social, technical, and ecological systems," *Journal of Artificial Societies and Social Simulation (JASSS)*, vol. 27, no. 1, 2024.

[24] J. Mitcham, "Agent-based simulation of police funding tradeoffs through the lens of legitimacy and hardship," *Journal of Artificial Societies and Social Simulation*, vol. 26, no. 3, 2023.

[25] F. Neves, P. Campos, and S. Silva, "Innovation and employment: an agent-based approach," *Journal of Artificial Societies and Social Simulation*, vol. 22, no. 1, 2019.

[26] J. M. Epstein and R. Axtell, *Growing artificial societies: social science from the bottom up*. Brookings Institution Press, 1996.

[27] E. Serrano and K. Satoh, "An agent-based model for exploring pension law and social security policies," in *New Frontiers in Artificial Intelligence: JSAI-isAI International Workshops, JURISIN, AI-Biz, LENLS, Kansei-AI*. Springer, 2020, pp. 50–63.

[28] M. Oremland and R. Laubenbacher, "Using difference equations to find optimal tax structures on the sugarscape," *Journal of Economic Interaction and Coordination*, vol. 9, pp. 233–253, 2014.

[29] A. Rahman, S. Setayeshi, and M. S. Zafarghandi, "Wealth adjustment using a synergy between communication, cooperation, and one-fifth of wealth variables in an artificial society," *AI & society*, vol. 24, pp. 151–164, 2009.

[30] P. Kurakin, "Technoscape: a multi-agent model of all-human global web," in *2022 15th International Conference Management of large-scale system development (MLSD)*. IEEE, 2022, pp. 1–5.

[31] J. A. Lasquety-Reyes, "Towards computer simulations of virtue ethics," *Open Philosophy*, vol. 2, no. 1, pp. 399–413, 2019.

[32] B. Sterner, "Agent-based computer simulation and ethics," *Metascience*, vol. 21, no. 2, pp. 403–407, 2012.

[33] D. A. Urbina and A. Ruiz-Villaverde, "A critical review of homo economicus from five approaches," *American Journal of Economics and Sociology*, vol. 78, no. 1, pp. 63–93, 2019.

[34] R. Pordes, D. Petravick, B. Kramer, D. Olson, M. Livny, A. Roy, P. Avery, K. Blackburn, T. Wenaus, F. Wurthwein *et al.*, "The open science grid," in *Journal of Physics: Conference Series*, vol. 78, 2007.

[35] D. Thain, T. Tannenbaum, and M. Livny, "Condor and the grid," *Grid computing: Making the global infrastructure a reality*, pp. 299–335, 2003.

[36] M. Albrecht, P. Donnelly, P. Bui, and D. Thain, "Makeflow: A portable abstraction for data intensive computing on clusters, clouds, and grids," in *Proceedings of the 1st ACM SIGMOD Workshop on Scalable Workflow Execution Engines and Technologies*, 2012, pp. 1–13.