Lecture 22: Artificial Agent Models: "Growing Artificial Society"

Agent-based models are computer game type simulations that use power of modern computers to examine how the interactions among artificial agents – computer code people/objects that act purposively in certain ways interact in digital environments to generate outcomes of interest. The models closely tied to cellular automata models. They are currently being used to model how people interact in different settings where they might spread the virus currently plaguing humanity and to assess social-distancing and other strategies to control and end the pandemic. They are a micro tool that can generate similar patterns to those in the differential equation based Susceptible-Infection-Recovery model of transmission with lots of details that could be useful. Here is an example from a German university

An agent-based policy laboratory for COVID-19 containment strategies

Ben Vermeulen*, Andreas Pyka**, Matthias Müller
* For inquiries about the conceptual or operational model, contact b.vermeulen@uni-hohenheim.de
** For inquiries into research collaboration as well as promotional and media usage, contact a.pyka@uni-hohenheim.de

Following the outbreak of COVID-19, governments are looking for a balance between containment on the one hand and retaining liberties of citizens and keeping the economy afloat on the other. Policy makers turn to academic modelers for policy guidance. Not surprisingly, academics are quick to present models also on the internet. For instance, few days ago, a simulation model of the COVID-19 diffusion process of The Washington Post was launched to explain the impact of 'social distancing'. In this model, the contamination process is driven by homogeneous agents roaming an undefined space and is thus more physics-based than social scientific. Moreover, the Neherlab, of the Biozentrum, University of Basel, recently launched a metapopulation model using a system of differential equations that essentially structure spatial features and social interactions.

In society, however, people generally infect each other in private and public spaces such as their own house, their office or school, in the supermarket, at mass gatherings such as sports events or concerts, etc. Moreover, certain people have central or gatekeeper positions in social networks and may thus be pivotal in spreading or containing COVID-19, while other people have peripheral positions and not even contract it. As such, people's agenda of activities, their social relationships, as well as the social setting of the locations matter a great deal. In contrast to the two models mentioned, agent-based models are particularly well-suited to study emerging socio-temporal patterns of real-world epidemics driven by the heterogeneity and autonomy of agents.

On this page, you find a first version of such an agent-based simulation model (with policy intervention features) in which agents are members of households, have a certain age and associated (basic) agenda determining where they spend time, and hence when with whom they interact – and potentially transmit the SARS-CoV-2 coronavirus. It is assumed that only when interaction takes place at particular locations, here: houses (orange), offices (purple), supermarket/ shop (red), school (pink), recreation & leisure (green), agents can infect one another. Moreover, there is a simple disease progression model akin to the SEIR metapopulation model specifying the infection transitions going from susceptible/ uninfected (green), exposed but latent/ non-infectious during incubation (light yellow), exposed and mildly presymptomatically infectious (yellow), infected and infectious (orange), severely ill (dark orange), critically ill (red), then either recovering and immune (purple) or deceased (blue).

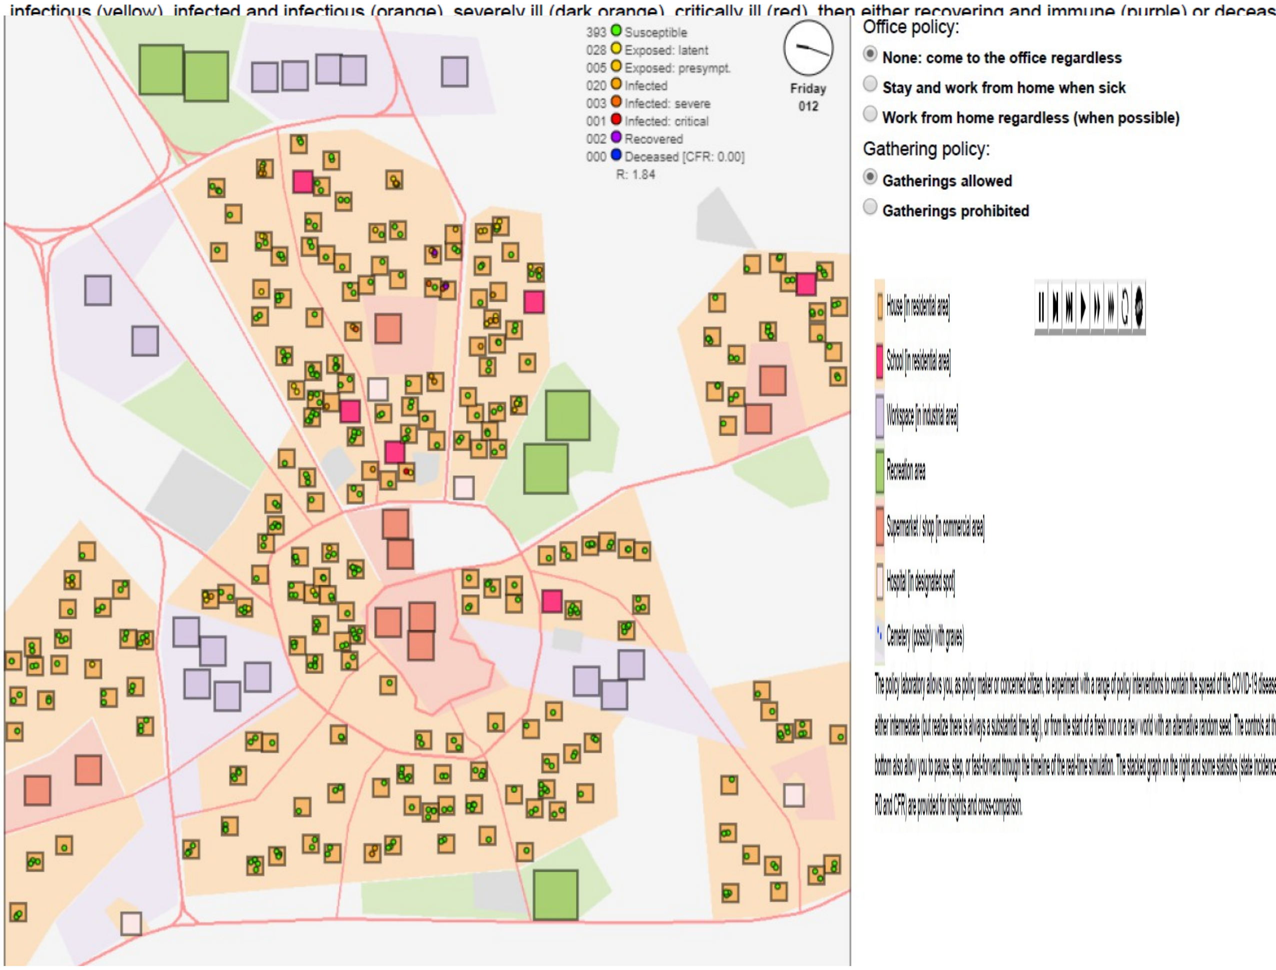

Leigh Tesfatsion (Iowa State) "Real-world economies are open-ended dynamic systems consisting of **heterogeneous interacting participants …** who strategically take into account the past actions and potential future actions of other participants. All participants are … **locally** constructive, meaning their actions at any given time must be based on their local states; and participant actions at any given time affect future local states ... these properties imply real-world economies are locally constructive sequential games... agent-based computational economics (ACE) … permits researchers to study economic systems from this point of view" SU Economics Working Paper 17022. March 17, 2017. Also, www.econ.iastate.edu/tesfatsi/ace.htm; http://www.econ.iastate.edu/classes/econ308/tesfatsion/ACETutorial.pdf

AA models live on LOCAL interactions of agents who adapt to market forces, using **inductive** rules – cellular automata writ large. They are opposite of systems dynamics models that live on expert knowledge of true system or macro models that live on modeler-imposed optimality and equilibrium conditions. For AA think micro-economics while for Systems Dynamics think macro-economics. The best AA models reveal surprising outcomes/ emergent phenomenon and regularities in data that local behavior does not obviously imply.

"...the software code is the model. But the software code … can equivalently be expressed as a system of discrete-time or discrete-event difference equations, starting from user-specified initial conditions. (Tesfatsion et al. (2017, sec6.3). "The role of the modeler is limited to the setting of initial agent states and to the non-perturbational observation, analysis, and reporting of model outcomes." She views them as a methodological approach, not a theory (though each model is a theory about some issue/process).

The models are "AI"-related. See https://en.wikipedia.org/wiki/Intelligent_agent

Point of this class is to introduce you to the models and show you the "architecture" that makes them work..

Examples of AA models in economics:

1) **Zero intelligence agent models** where competitive outcomes arise from non-rational/imperfect equilibrium. Simplest case is where firms randomly choose a strategy but market interactions produce something akin to competitive equilibrium. Random choice → competitive outcome because surviving firms must choose what consumers want. **Minority game** that produces "close to coordinated solution" based on lookup table rules for agents with limited induction of common information. Key is that uncoordinated actions → close to coordinated outcome.

2)**Schelling cellular automata** models where local preferences –> excessive segregation/outcomes no one wants.

3)**Sugarscape model** where agents follow simple rules of discovery and trade –> Pareto distribution of incomes

4) Random errors –> Normal distribution/Bell-shaped curve. Other stable distributions – Cauchy-Levy .

Modeler must specify artificial Agent in terms of computer code that specifies BEHAVIOR that adapts to environment. The agent has

Sensors – so it "sees" objects a certain distance on grid; Identifies the objects via some tag (Red = good)
Internal representations to assess the value of objects– **utility function** – PURPOSIVE
Ability to make inferences – modifying internal data representations. Code finds out that red is good
Effectors that modify environment (behavior): move; **change state**; communicate/trade with other agents

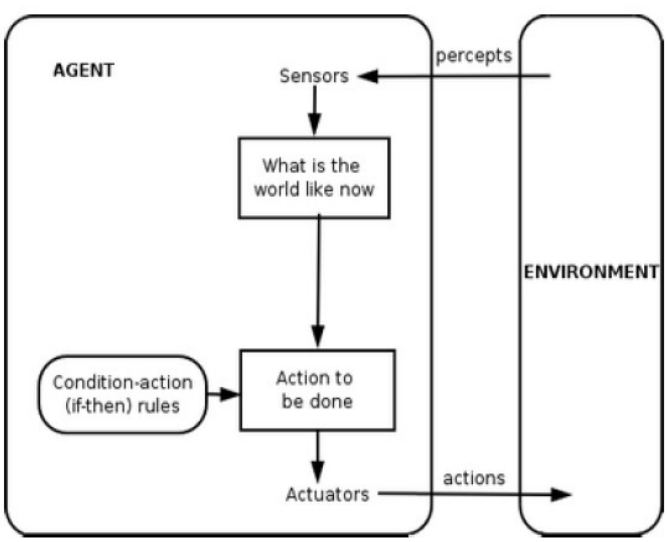

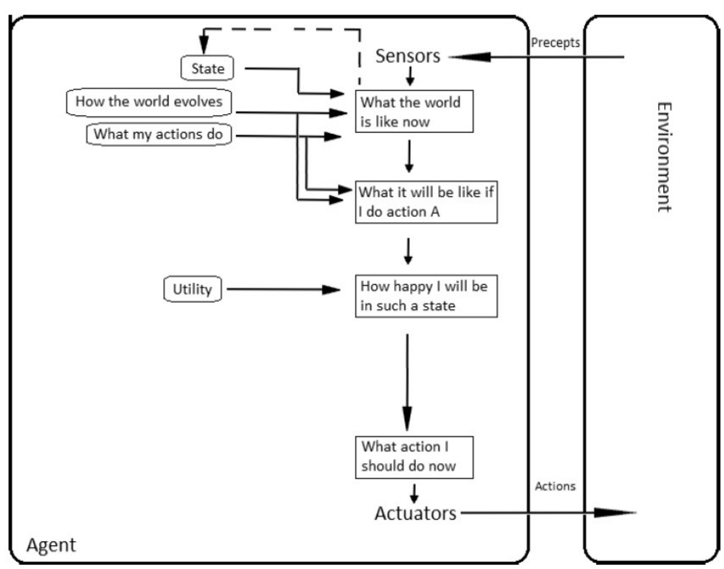

AA simulation models are THEORY based on heterogeneous agents and CA style rules that evolve over time.

Binmore criticizes the approach: **simulation is no substitute for thinking** – ie you won't find the normal distribution from the simulation – just some bell-shaped curve. But as computers allow more and more complicated simulations, maybe simulation based on local knowledge can beat out economists "thinking" (rigid rational expectations model that assumes everyone (but the analyst) knows the structure of the world. etc). A substitute for thinking? AlphaZero!

Much hype for AA as Complex Adaptive Systems (CAS). A CAS is a system with local agents, so that coherent behavior arises from competition/cooperation among the agents. System LEARNS or EVOLVES by revision and recombining the building blocks. Genetic algorithm or some other learning/change behavior/evolve rule will be important. Check http://socdynamics.org/id4.html. For learning, AI. AlphaZero with a data-streaming twist to decide when "game" has changed. Give the algorithm way to shift from Chess to Go to Poker???

Epstein-Axtell: "most effective way to alter collective patterns of behavior may be from the bottom up, by modifying local rules". Holland (inventor of Genetic Algorithm): "CAS has lever points, wherein small amounts of input produce large directed changes. It would be easier to discover these lever points if we can uncover general principles that govern cas dynamics " (Holland, 39-40).

Model 1 Zero Intelligent Agents Produce Efficient Outcomes (Gode/Sundar)

Builds on insight that key economic relations depend critically on exit and entry rather than rational optimizing behavior. We get downward sloping demand curves even when people pick quantities randomly. Why? Because budget constraint – distribution of opportunities – forces downward slope. When the price of a commodity rises a lot you cannot afford as much of it as you could at lower price. Consider zero-intelligent consumers who at each each randomly call out how much to buy, so they are not making any rational calculation. Budget constraint means that they can only buy amounts to the left of the budget "constraint curve". Take the mean values bought at the different prices and you get a downward sloping demand curve.

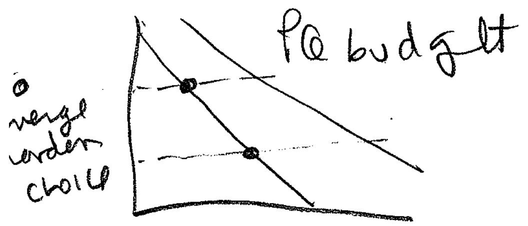

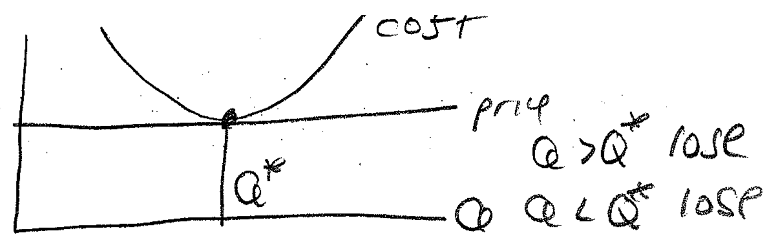

B- Brainless Producer Firms randomly make decisions. The firm that gets the right output survives. Others go out of business. In fact survival rates of firms are modest; 30-40% die in years 1-2 and the half-life is 7 years for a firm

Gode-Sunder show that "Zero Intelligence (ZI) programs" that submit random bids and offers but only accept bids/ offers that meet a budget constraint so they sell above the production cost and buy at less than the value → long term efficient outcomes. The model is a **Double auction** with 6 buyers and 6 sellers. Each seller has a single unit; each buyer wishes to purchase one unit. Buyers have value Vi to a unit. Sellers have cost Ci. Buyer gets Vi-P. Seller gets P-Ci. This generates supply and demand. In human experiment people call out values and reach price. They make deals. The bids order from highest to lowest into a demand curve and offers order from lowest to highest into a supply curve. P is a weighted average of the n-th and n + 1-th bids. Buyers whose bids are above and sellers whose bids are below, this market price, buy or sell respectively one unit at the prevailing price.

Double auctions in lab experiments with people find simple good markets clear very fast; more complicated goods markets may not clear at all. HERE ARE PEOPLE RESULTS

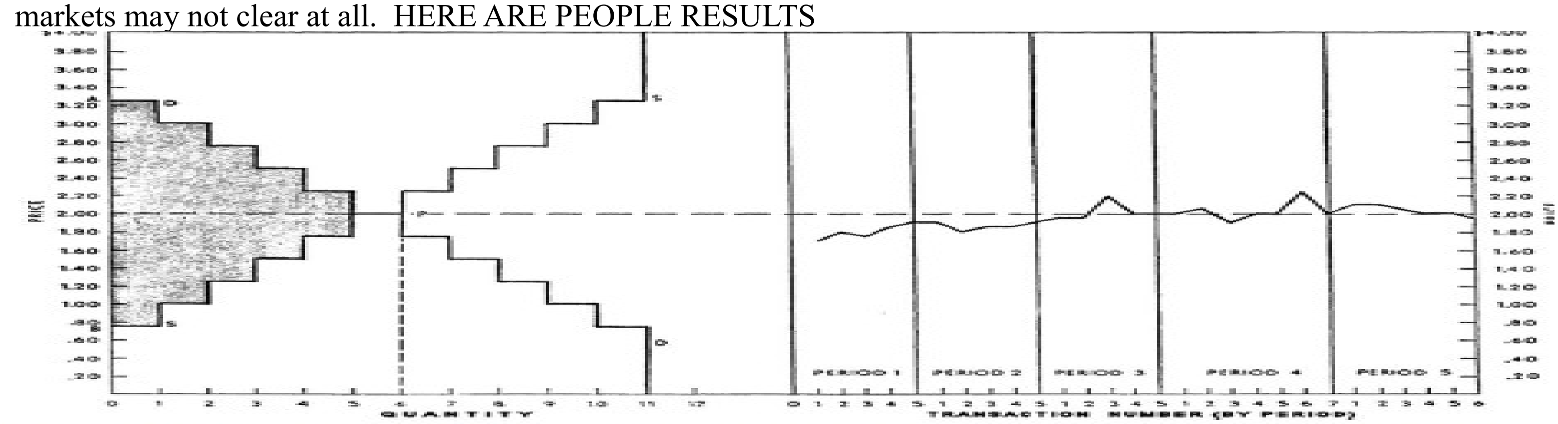

...re 1. Values and costs induced in an experimental double auction design (left panel) and the path of prices achieved by human subjects (right panel). S... Smith (1962, Chart 1).

Here are people results vs zero-intelligent agent results **WITHOUT vs WITH** BUDGET CONSTRAINTS:

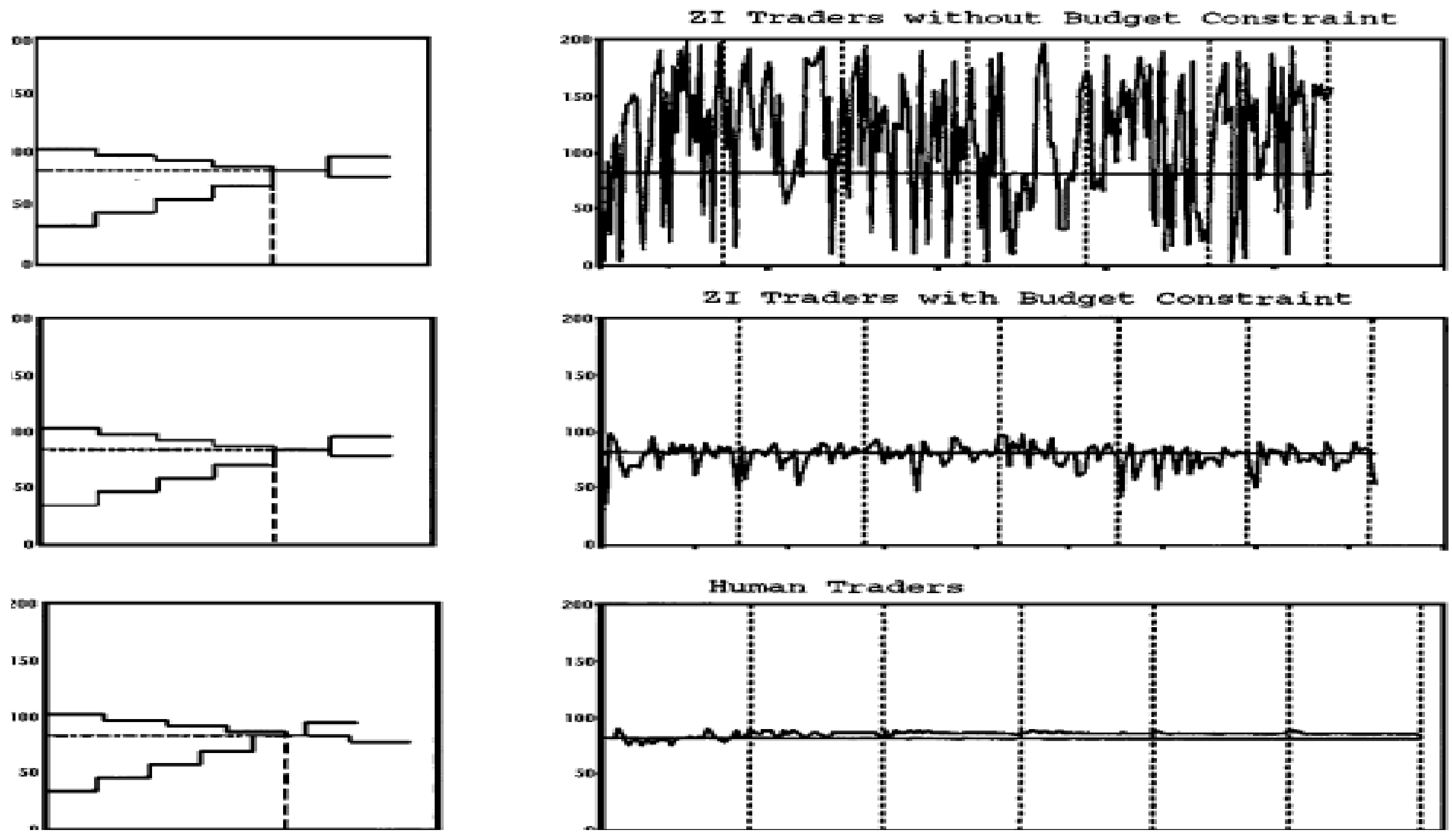

Criterion for Efficiency: % of Consumer Surplus + Producer Surplus effectuated
Does this mean that market institutions determine results and human rationality/cognition are unimportant? Just give me constraints and a market framework and I will tell what the results will be.

MEAN EFFICIENCY OF MARKETS IN FIGURE 7

Traders	Market 1	Market 2	Market 3	Market 4	Market 5
ZI-U	90.0	90.0	76.7	48.8	86.0
ZI-C	99.9	99.2	99.0	98.2	97.1
Human	99.7	99.1	100.0	99.1	90.2

Zero Intelligence (ZI) vs "Realistic Market Agent"

ZI - accurately model market mechanism and constraints; assume individuals have no strategy and behave at random to identify the effect of market rules. Anything else must require the interaction of strategies of market participants. Goal is to separate the effect of market mechanism from strategy and optimizing behavior.

Weakness: ZI traders do not learn from their experiences, and as a consequence they poorly replicate experimental results when a market situation are repeated or expert traders are employed. If the rules of the mechanism are not accurately captured, the dynamics of the model may be misleading. They do not make "insiders' deals" that are often the key money-makers nor predict any changes in rules

Duffy "The ZI approach is perhaps best suited to competitive environments, where individuals are atomistic and... institutional features together with constraints on unprofitable trades will largely dictate the behavior that emerges. In environments where agents have some strategic power, so that beliefs about the behavior of others become important, the ZI approach is less likely to be a useful modeling strategy."

2. EXAMPLE: Sugarscape as an Archetype system

Heterogeneous agents with different abilities and "tastes" for sugar and spice look for food on a lattice:

1) The lattice gives different attributes to locales – think computer game with different colors to represent water, mountains, plains. The locales allow the AAs to do some things and not others.

2) AA operate on this space and interact with other AAs according to certain rules. The AAs have utility functions: $\log W_1 = m_1/(m_1 + m_2)$; $\log W_2 = m_2/(m_1 + m_2)$, where ms measure value of the items, sugar or spice.

The agents have vision and move according to specified rules.

Figure II-3. Agent Vision

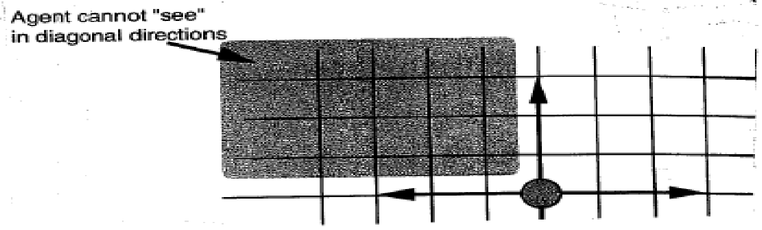

Agent movement rule **M**:
- Look out as far as vision permits in the four principal latti directions and identify the unoccupied site(s) having the mo sugar;[8]
- If the greatest sugar value appears on multiple sites th select the nearest one;[9]
- Move to this site;[10]
- Collect all the sugar at this new position.

Ahe rules produce a concentration of agents at the spaces with the desired/needed resources and skewed income aistribution (how much is due to luck and how much to "ability"/process" – how much could you tax away??

NET LOGO: DOWNLOADABLE AND EASY FOR PAPERS http://ccl.northwestern.edu/netlogo/index.shtml

Animation II-2. Societal Evolution under Rules ({G_1}, {**M**}) from a Random Initial Distribution of Agents

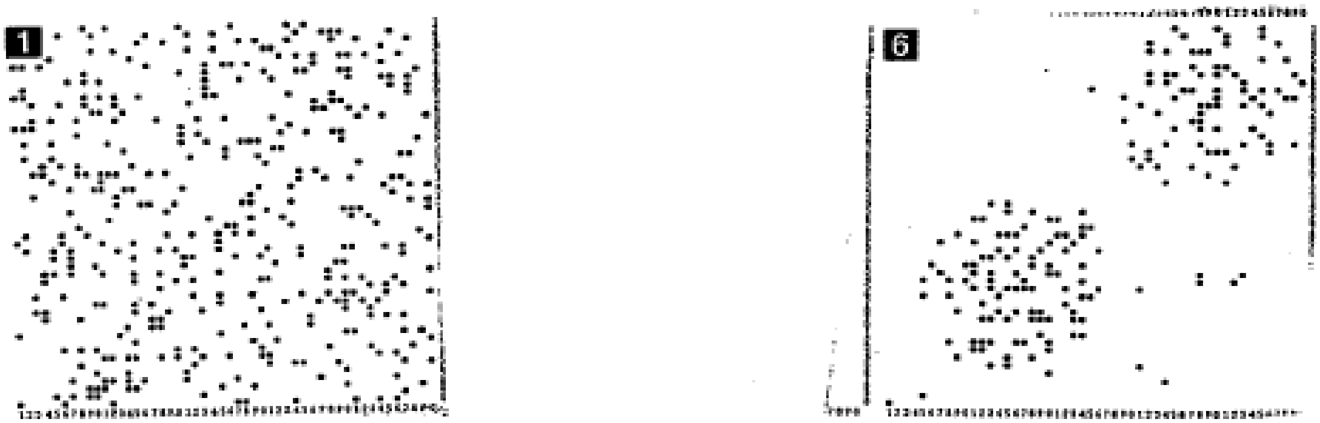

And a highly skewed income distribution.

Animation II-3. Wealth Histogram Evolution under Rules ({G_1}, {**M**, $R_{[60,100]}$}) from a Random Initial Distribution of Agents

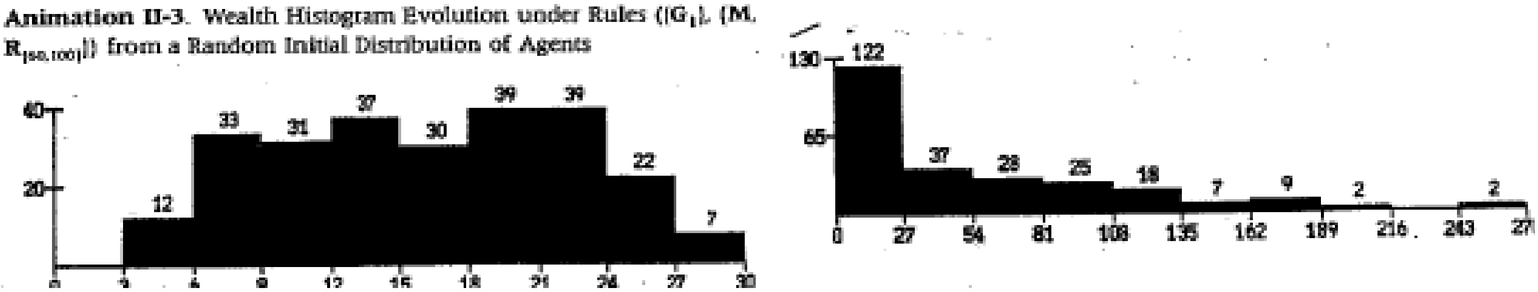

Addition of LOCAL trade –based on the lattice. You have spice, neighbor has sugar. You "bargain" over price and trade according to the trade rule and move according to a new move rule: get anything you can since you can trade it.

Multicommodity agent movement rule **M**:[7]
- Look out as far as vision permits in each of the four lattice directions, north, south, east, and west;
- Considering only unoccupied lattice positions, find the nearest position producing maximum welfare;
- Move to the new position;
- Collect all the resources at that location.

Agent trade rule **T**:
- Agent and neighbor compute their *MRS*s; if these are equal then end, else continue;
- The direction of exchange is as follows: spice flows from the agent with the higher *MRS* to the agent with the lower *MRS* while sugar goes in the opposite direction;
- The geometric mean of the two *MRS*s is calculated—this will serve as the price, p;
- The quantities to be exchanged are as follows: if $p > 1$ then p units of spice for 1 unit of sugar; if $p < 1$ then $1/p$ units of sugar for 1 unit of spice;
- If this trade will (a) make both agents better off (increases the welfare of both agents), and (b) not cause the agents' *MRS*s to cross over one another, then the trade is made and return to start, else end.

EPSTEIN PUP https://en.wikipedia.org/wiki/Comparison_of_agent-based_modeling_software lists http://ccl.northwestern.edu/netlogo/index.shtml Josh Epstein Agent_ZeroToward Neurocognitive Foundations for Generative Social Science (PUP 2014),

Wallace, R., Geller, A., Ogawa, V.A., Eds. (2015). Assessing the Use of Agent-Based Models for Tobacco Regulation. Institute of Medicine of the National Academies. National Academies Press:Washington, D.C.

Model III --**NASDAQ tick model Sixteenths or Pennies?** Darley, Outkin, Plate, Gao**.Simulates Nasdaq**

Tick Fight: Argument between NASDAQ and SEC over size of ticks – unit of monetary transaction and spread. SEC wanted Nasdaq to move from quoting prices in sixteenths of a dollar to decimal prices, of one hundredth of a dollar. **This reduced the Nasdaq "tick size" from $0.0625 to $0.01. SEC argued that smaller tick size of transactions would encourage competition and reduce the spread in the share market.** On June 24, 2014 the SEC issued an order directing that the national securities exchanges and FINRA (collectively, "SROs") to act jointly to develop and file an NMS plan with the SEC to implement a pilot program to study the impact of wider quoting and trading increments on trading of certain small capitalization stocks. On May 6 2015 the SEC approved **Tick Pilot NMS Plan**

/

Last time Nasdaq changed its tick-size: "On March 13, 1997, the American Stock Exchange (AMEX) and Nasdaq announced plans to reduce their minimum tick size from US$0.125 to US$0.0625. Since that date, the New York Stock Exchange (NYSE), AMEX, and Nasdaq have all reduced their minimum tick size." -->change of strategies.

Nasdaq commissions Darley & Outkin to build a simulation model to evaluate effects of changing tick-size to $0.01 before SEC final decision. They build "a model that represents a highly realistic picture…of a dealer-mediated market, with … many features of real-world markets" and analyze consequences of regulatory and structural changes to the market (ie minimum tick size). "While building the simulated model … we interacted extensively with many market participants: market-makers, brokers, traders, large investors, etc. We found this interaction invaluable – as a source of information .., on often subtle details of market operations, . One conversation with a market maker still stays clear in our minds. He was supportive, but skeptical... his skepticism lay in this question: how one can model the fear and greed often ruling market behaviour? This is a valid point …understanding of … individual and mass psychology is lacking."

They design AAs for market makers (dealers) and investors, whose interactions produce price discovery and determine the markets dynamics; investigate behaviors as AAs learn to perform profitably. Examine whether the "model, at least in a stylized fashion, is able to replicate some of the observed features of real-world markets."

Claims: "model is robust in that the simulated market, exchange, investors, and dealers perform realistically under a wide variety of conditions … market dynamics produced by the model have the same qualitative properties as those observed in real markets... pertaining to volatility, liquidity, spread sizes, and spread clustering… and thus provides a test bed in which to investigate the effects of changes in market rules and conditions …

the simulation suggests that a reduction in the market tick size ... from $1/16 to a penny c**an reduce the market's ability to perform price discovery,** particularly when **parasitic strategies** such as SOES bandits and day traders (who make a small transaction to change price and then make money by following with larger transaction at new price) are present in the market."

Our results are "not a consequence of a uniquely identifiable feature of the model, or of the actions of certain market participants. Rather, **they result from a relatively complex set of interactions** of market makers, investors, market rules, and market infrastructure. Thus, even in a relatively simple setting, we can observe unintended consequences of the market s design — for instance, spread clustering, which occurs in our simulation. Their model said, reducing tick size would be bad because it would increase # of **SOES bandits** and thus create more volatility:

" Introducing finer tick sizes for US equities could put a strain on long-only buy-side desks and give an additional edge to high-frequency firms, according to senior traders. US exchanges, including the New York Stock Exchange and Nasdaq, are reportedly lobbying the US Securities and Exchange Commission (SEC) to permit them to accept quotes in increments finer than a cent. Currently, exchangescan accept quotes in increments of a hundredth of a cent for shares trading below $1. But the exchanges are said to favour reducing the tick size to a tenth of a cent for stocks trading at or above $1 to allow them to compete more effectively with non-displayed trading venues.

Outkin (2012, http://computationalsocialscience.org/wpcontent/uploads/2012/09/Outkin2012.pdf) says most predictions of the model came out. But Congress spoke

House Pressures SEC to Widen Stock "Tick Prices" WSJ 13 February 2014

Lawmakers are ratcheting up pressure on U.S. securities regulators to allow stock prices of smaller companies to be traded in increments of more than one cent, a proposal some argue will boost the ability of small companies to raise money from investors. The House late Tuesday voted 412-4 to pass a bill, sponsored by Reps. John Carney (D., Del.) and Sean Duffy (R., Wis.), to require the Securities and Exchange Commission to create a pilot program allowing companies with stock market values of less than $750 million to trade in "tick sizes" of five cents or 10 cents. The testing phase would last for up to five years.

In response SEC sets up experiment with The Pilot Securities were divided into one control group and three test groups. Each test group contains approximately 400 Pilot Securities and the remaining Pilot Securities are in the control group. The groups are defined as follows:

- Test Group One quoting in $0.05 per share increments but trade at the current price increments
- Test Group Two quoting in $0.05 per share increments like those in TG1, but are trading in $0.05 per share increments, but permits executions that are the (1) midpoint between the best bid and best protected offer, (2) retail investor orders with price improvement of at least $0.005 per share.
- Test Group Three quoting in $0.05 per share increments and trading in $0.05 per share increments. Pilot Securities are also subject to a Trade-at Prohibition, which generally prevents price matching by a trading center that is not displaying the best price unless an exception applies.
- **Control group** continue to quote and trade at the current tick size increment of $0.01 per share.

The exchanges and FINRA will conduct a joint assessment of the Tick Size Pilot's impact based on the first year's data and will provide the assessment to the Commission and make it publicly available no later than April 2018.

What did the experiment show?

Tick Size Pilot Plan and Market Quality   January 28, 2019   https://www.sec.gov/files/dera_wp_tick_size-market_quality.pdf   Overall, we find that on average, relative to stocks in the Control Group, market quality deteriorates for stocks in the Test Groups. Specifically, we find that stocks in the Test Groups experience an increase in spreads and volatility and a decrease in price efficiency, relative to stocks in the Control Group.

Tick Size Pilot Plan Threshold Analysis March 6, 2019 https://www.sec.gov/files/dera_wp_ticksize-thresholdanalysis.pdf We do not find clusters of stocks which systematically experienced improvements in market quality. Few stocks experience improvements in market quality and these stocks are not clearly identifiable based on pre-Pilot characteristics.   Based upon visual inspection, we are unable to identify clear thresholds for several stock characteristics above or below which stocks experience an improvement in market quality.
So the simulation model said big problems if reduce tick size. The SEC experiment says that going from 0.01 to 0.05 would worsen market and not obvious if there is threshold where raising tick would be beneficial. Depth measures how big an order is needed to change the price.   - The National Best Bid and Offer (NBBO) is an (SEC) regulation requiring brokers to trade at the best available ask (lowest) price and the best available bid (highest) price when buying and selling securities for customers.

"Economic reasoning and artificial intelligence" Parkes and Wellman, (SEAS) SCIENCE 17 JULY 2015 • VOL 349 ISSUE 6245 aka **Good buy People, with all your cognitive biases; Hello MACHINA ECONOMICUS**

"AI strives to build rational agents capable of perceiving the world around them and taking actions to advance specified goals" – a synthetic **machina economicus,** based on purely rational optimizing decisions with none of the self-defeating cognitive biases that impair human decisions. Paper discusses challenges in designing AIs that can reason effectively in economic contexts. Most references are to Computer Science --

1.Growth of AI agents making decisions, setting prices, making bids, etc . Current: automated trading algorithms estimated to be responsible for more than 70% of trades on U.S. stock markets; Amazon book pricing, …

2."Without offering any judgment on the question of how well rationality theories capture essential human behavior, we note the irony in the prospect that social science theories may turn out to apply with greater fidelity to nonhuman agent behavior." But of course, AI reasoning based on rationality, game theory assumptions that other players/AI play strategically, etc is the world of economic theory – mechanism design, etc.

3.Search engine auctions have supported ad bidding algorithms in ways that fit auction/equilibrium models. Reputation systems set up for reporting on buyers and sellers. Continual need to seek incentive-compatible rules so that AI does not exploit its position for narrow gains that reduce global output. And must consider cost of computation.

4."As AI advances, we are confident that economic reasoning will continue to have an important role in the design of single-agent and multi-agent AIs, and we have argued that, as conomies of AIs continue to emerge, there will need to be a new science to understand how to design these systems. These AIs will no doubt exert strong forces on the economy and broader society; understanding the effect and extent of this will shape the research agendas of both AI and economics in years to come.

Table 1: Example Covid-19 Forecast and Projection Models for the U.S.			
Model and Organization(s) Responsible	**Primary Approach**	**Outcomes Estimated and Timeframe**	**Selected Model Findings/Notes**
Imperial College "non-pharmaceutical intervention" (NPI) Model	SEIR	Projected U.S. cases, deaths across a range of different mitigation and suppression scenarios, over the next year (to April 2021).	Projected 2.2 million U.S. deaths might occur in an "unmitigated" scenario
Institute for Health Metrics and Evaluation (IHME) Covid-19 Model	Curve-fitting/ extrapolation	Forecasts number of hospitalizations and deaths in the U.S. and by state, along with the timing of in the peak of hospitalizations and deaths, through August 2020.	Initially, the model forecast 81,000 deaths in the US by July. Results are updated daily, and as of Apr 12, that deaths estimate has been revised downward, to 61,545 by August 4.
Covid-19 Model from Northeastern University, Fogarty International Center, Fred Hutchison Cancer Center, University of Florida and others	Agent-based	Projects cases and deaths in the U.S. and by state, under no mitigation vs. "stay-at-home" scenario, through April 30, 2020.	As of April 4, model projected U.S. deaths would peak on April 8, and there would be approximately 52,575 COVID-19 deaths (range: 35,381 to 88,269) by April 30, 2020
Columbia University Severe Covid-19 Risk Model (& mapping tool)	SEIR	Provides projections on number of severe cases, hospitalizations, critical care, ICU use, and deaths under different social distancing scenarios, for 3-week and 6-week periods starting April 2.	In different regions of the U.S. anywhere from 33,986 and 185,192 deaths could be averted through social distancing.
Los Alamos National Laboratory Confirmed and Forecasted Case Data Model	Curve-Fitting/ Extrapolation	Forecasts cases and deaths by U.S. state using assumptions about the growth rate in cases and deaths and the presence of social distancing interventions through May 20.	As an example, model best guess forecast for California as of April 8 is that there would be 138,100 cases and 4,082 deaths.
University of Pennsylvania Covid-19 Hospital Impact Model for Epidemics (CHIME)	SIR	Model allows users to set inputs and assumptions, then provides forecasts on expected number of hospitalizations, ICU bed demand, ventilator demand, and number of days these demands would exceed capacity at hospitals in a given area based on those inputs, over the next three months.	Using inputs for three University of Pennsylvania Health System hospitals, the model projected best- and worst-case scenarios for total hospital bed capacity needed would reach 3131 – 12,650, including 338 – 1,608 ICU beds and 118 to 599 ventilators.