The Market Algorithm as Distributed Genetic Algorithm

Zachary Johnson
North Dakota State University
Fargo, ND, United States
zachary.e.johnson@ndsu.edu

James Lee Caton Jr
Challey Institute for Global Growth and Innovation
North Dakota State University
Fargo, ND, United States
james.caton@ndsu.edu

Dragan Miljkovic
North Dakota State University
Fargo, ND, United States
Dragan.miljkovic@ndsu.edu

Veeshan Rayamajhee
Challey Institute for Global Growth and Innovation
North Dakota State University
Fargo, ND, United States
Veeshan.rayamajhee@ndsu.edu

Keywords: Agent-based Computational Economics; Artificial Intelligence; Entrepreneurship; Evolutionary Economics; Macroeconomics;

The Market Algorithm as Distributed Genetic Algorithm[1,2,3,4]

Abstract

Portraying the market as a genetic algorithm requires that we define where in the market process 1) mutation and 2) crossover (partial duplication) occur. Attributes of an agent are mutated at a modest rate and the attributes of the best performing agents are mixed through crossover, giving a child agent attributes of its parents. In the marketplace, entrepreneurs learn from one another and don't necessarily know which strategies produce globally optimal results. To isolate the effectiveness of market competition on agent rationality, it is critical that we use simple heuristic agents as an analytical baseline. Our model contributes to the evolutionary literature by analyzing outcomes through the lens of economic theory. While there is only capital in our model, we evaluate economic efficiency using a simple macroeconomic model of capital productivity where improvements in the sustainable level of capital indicate improvements in the average productive efficiency of capital. We compare performance of neoclassically rational agents to performance of agents who employ a complex of simple heuristics refined via the competitive market process. We find that competitive pressures lead to comparable results between these two classes of agents, with heuristic agents being only modestly less productive than neoclassically rational agents. We interpret the model as an expansive NK landscape where fitness is indicated by wealth per capital and the level of capital.

Introduction

It is well known that competitive markets process information. This was the insight behind the efficient markets hypothesis (Muth 1961; Fama 1970). It was this insight that allowed Armen Alchian to discover the element used for fuel in the hydrogen bomb (Newhard 2014). And it was this insight that defined much of F.A. Hayek's work concerning market processes (1933; 1935a; 1935b; 1937; 1940; 1945; 1948; 1967; 1968ş see also Caldwell 2008). Bowles, Kirman, and Sethi (2017) track Hayek's argument, voicing both appreciation and skepticism. Herding behavior, they note, is both important for the functioning of markets and also integral to speculation that can yield "considerable nonfundamental volatility" in prices.[5] The authors, concerned with disequilibrium, recommend agent-based modeling as a means of

---

[1] An incomplete list of those whose input was helpful with regard to the article include Jeremy Straub and Cameron Harwick. Thank you to the economics department at Universidad San Francisco de Quito for inviting a very early draft of this project.

[2] Funding Declarations:

The authors express their gratitude to the funding provided to support this study from National Science Foundation (NSF) EPSCoR RII Track-2 Program under the grant number OIA-2119691. The findings and opinions expressed in this article are those of the authors only and do not necessarily reflect the views of the sponsors.

This work used resources of the Center for Computationally Assisted Science and Technology (CCAST) at North Dakota State University, which were made possible in part by NSF MRI Award No. 2019077

[3] The authors have no competing interests as defined by Springer, or other interests that might be perceived to influence the results and/or discussion reported in this paper.

[4] Data Declaration:
Data is available upon request. Several gigabytes of data have been generated and can be provided along with script for model and analysis and presentation of results. Model can be found at:
https://github.com/zacharyejohnson/Academics/tree/main/Sugarscape

[5] In their discussion, the authors refer to Hayek's example of tin.

better understanding market processes (for example, Bylund 2015; Holian and Newell 2017; Keyhani 2016, 2019).

In “Friedrich Hayek and the Market Algorithm”, Bowles, Kirman, and Sethi (2017) provide some guide posts for thinking about the market algorithmically, especially in their emphasis on the information role of prices, economic disequilibrium, and herding. The authors recognize that there is no canonical structure for agent-based modeling to guide those who would like to use the tool to develop insights about market processes. Even a half-century after Schelling’s segregation model, agent-based modeling still operates outside of the mainstream of economic theory[6].

Entrepreneurs that drive market processes are not represented in the neoclassical equilibrium theory of Jevons and Walras. William Baumol lamented this fact, noting that “[t]he Prince of Denmark has been expunged from the discussion of Hamlet (Baumol 1968).”[7] Economics, as the science of decision-making, recognizes that action is ends oriented. Once ends have been determined, agents must evaluate which resources best serve as means to promote manifestation of desired ends. Agent actions evidence their valuations to the observer. Action reveals preference. This is true of the neoclassical agent who adjusts his bundles according to the calculus of optimization as much as it is true for the man or woman who has never been introduced to economic theory and whose ends are multitudinous. These attributes of individual decision-making aggregate to produce market outcomes often similar, but at times at variance, to the equilibrium state predicted by textbook neoclassical theory.

In developing a systematic description of the market algorithm, we develop Hayek’s forward looking concern that “tendencies to make economic theory more and more formal . . . have not yet been carried far enough to complete the isolation of this branch of logic and to restore to its rightful place the investigation of causal processes, using formal economic theory as a tool in the same way as mathematics” (1937, 35). Hayek sought a formal theoretical structure capable of meaningfully integrating methodological individualism as a guiding principle. Such a formal apparatus could confront prescient questions concerning the significance of the development, distribution, and coordination of knowledge. This includes decentralized price formation consistent with Adam Smith’s description of the market as “a system of extreme decentralization” (Coase 1992, 713).

Economic theory cannot be wholly contained within the neoclassical paradigm as descended from Jevons and Walras. When pushed, most economists would not accept that humans inherently act or ought to act as described by the neoclassical model.[8] But nearly all of us would agree that the model describes the outcome that competitive markets supported by robust institutions tend to generate. The process that generates these outcomes, however, does not operate with the machine-like accuracy of the neoclassical model. The neoclassical model often produces the correct inference. However, its representation of agent rationality is intended more for the convenience of marginal analysis than faithful representation, and it is incapable of modeling market processes that generate the outcomes it predicts. The detailed theoretical form from the Mengerian line provides appropriate foundations for

---

[6] In Schelling’s segregation model, agents are initially distributed randomly on a grid, moving only if their preference for neighbor similarity is not met. The model shows that even a preference that one in 3 neighbors is of the same race can lead to segregated neighborhoods.

[7] For a presentation of neoclassical theory that makes use of the entrepreneur to bring meaning to the theory, see Kirzner (1963).

[8] For a useful example of a neoclassical description of human behavior see Stigler and Becker (1977)

description of the market algorithm (Menger 1871). Our framework allows for fine grained investigation of this process.

Precedent for an Agent-based Framework

Bowles, Kirman, and Sethi "find considerable lasting value in Hayek's economic analysis while nonetheless questioning the connection of this analysis to his political philosophy (2017, 217)." As with any theorists, some claims from Hayek are foundational to economic analysis and some conclusions require faith in additional assumptions either not clearly defined or not required to be adopted in analysis. Using Hayek's example of tin, these same authors argue that the information content of prices can be distorted by herding mentality amongst investors.[9] We do not wish here to present a counterargument, though we think their point is worthy of discussion and should ultimately be considered within the academic dialectic that is supposed to refine philosophical and technical knowledge (Polanyi 1951; Mokyr 2017). Rather, we wish here to argue that holders of diverse ideologies can – and perhaps should – seek an agreed upon framework of understanding to facilitate debate. The only other solution would be to discard the value of dialectic in regards to its ability to uncover truth and instead presume that theorists of different schools are inherently siloed by differences in theoretical approaches.[10] One might choose to believe the latter, but the tradition of scholarly discourse traced back through the fountainheads of our discipline – most recently through figures like Keynes or Hayek, Friedman or Samuelson, Coase or Simon – show an appreciation for the need to explain new ideas with reference to ideas that preceded those new ideas and in a manner that is intended to appeal across our profession.

There is a body of economic theory concerning which most economists tend to agree. There is an overwhelming consensus that, in the long-run, markets tend to guide resources toward their most valued uses, as long as the existing institutional framework keeps transaction costs sufficiently low (Coase 1937; Coase 1960; Williamson 1985). You might as Keynes (1936), accept this premise and attempt to demonstrate exceptions to the rule that complicate our understanding of market processes. Keynes proposed that the market for money was different from markets for other consumption goods and services. For him, financial markets were subject to herd mentality that drove harmful oscillations and tended to promote macroeconomic disequilibrium that prevented the macroeconomy from reaching an optimal equilibrium. Similarly, he considered demand to hold money during a downturn to be implacable except by debt-financed fiscal expansion. While we do not develop money in this emergent model, we leverage herding as part of a rudimentary form of expectation formation.

Agent-based modeling presents an opportunity to demonstrate the significance of economic principles and the nature of exceptions to textbook neoclassical theory. In particular, an evolutionary paradigm like the one prized by Hayek (1962; 1967; 1968) is well-suited to elaborate an evolutionary economic theory.[11] As Alchian (1950) pointed out, the postulate of agent rationality that was growing in popularity

[9] Hayek indicates that the price of tin reflects local perceptions of actors intending to use this resource as a productive input. As the price of tin responds to newly formed entrepreneurial plans, other users of tin will adjust the intensity of their demand for tin by substituting away from or toward tin, depending on whether the price of tin has increased or fallen.
[10] This is, for example, the approach taken by Marxian theorists Resnick and Wolff (1987).
[11] See also Vriend (2002).

at the time can be approximated by a mixture of experimentation, herding, and evolutionary selection.[12] At the time, economists lacked the appropriate toolset for demonstrating systems level effects on agent rationality, but this has not been the case for several decades now. We can effectively model Alchian's proposal to show how outcomes in the market often approximate the outcome that would be expected if agents were perfectly knowledgeable with regard to objective data in the environment *and* the subjective data of other agents.

While it is the case that agent-based modeling has much to contribute to the evolutionary paradigm, the best approach is yet to be adequately demonstrated. As Bowles et al. (2017) point out, there exists a "large and heterogeneous collection of models . . . under the umbrella of *agent-based computational economics*" that are capable of demonstrating economic disequilibria, but "there does not yet exist a canonical agent-based framework" to help coordinate these efforts (p. 224-225). Nelson and Winter note that the evolutionary paradigm requires "*behavioral continuity*" and that "the natural selection argument is based on *profit-induced growth*; that is, successful firms earn profits and expand" (2002, 27). Profit is a function of the fitness of the array of strategies developed within a firm (Winter and Nelson 1982). We identify an analytical baseline for evolutionary agent-based models that ties performance of an agent-based economy interpreted through the emergence of macroeconomics aggregates. This allows us to avoid the heroic assumptions of standard macroeconomic modeling (see Dosi and Roventini 2019).

In what follows, we propose features of agent-based models of market process. These support both systems level features and individual processes not captured by textbook neoclassical models (See for example, Arrow and Hahn 1971). In presenting this framework we contribute to the same scholarly conversation to which the economists employing the more popular equilibrium framework have contributed. We show how subjective and, often, myopic agents should ever develop expectations that, on the whole, tend to correspond to existing conditions (Hayek 1937, 29). Our agents form plans sequentially, along the lines envisioned by Radner, who, like Hayek, was comfortable with recognizing agent myopia (Glasner 2022). Our heuristic agents develop simple strategies learned over the course of production and bargaining. We compare the performance of agents that employ neoclassical rationality with CES utility functions to agents who employ simple heuristics. Following Alchian, knowledge of heuristics depends upon experimentation, (imperfect) herding, and evolutionary selection. Our simulation of a self-ordering, competitive market is analogous to a distributed genetic algorithm.

We consider structure for equitably comparing exchange in various modeling contexts. We then present our iteration of the Sugarscape model, which includes agents whose decisions are guided by heuristics as well as those guided by the calculus of optimization. We find that these optimizing agents perform only modestly better than heuristic agents when progress is compared over computing time. Experimentation under the duress of competition plays a critical role in refining strategies. Further, we find that increased experimentation requires greater reliance on herding strategies in order to swiftly discard failed decision rules in favor of successful rules. Aggregate capital dynamics follow similar patterns for both types of agents, with savings per capita converging to zero as the capital stock grows.

Time, Exchange, and Information in an Agent-based Model

---

[12] See also Becker (1962) and Kirzner (1962).

For some time now, agent-based models have shown how a static equilibrium might be approximated among exchanging agents with fixed preferences. Gode and Sunder's (1993) path-breaking model demonstrated that zero-intelligence traders can, by bumbling through exchange processes, approximate the equilibrium price and quantity for a stock that would be predicted by supply and demand curves implied by the preferences of trading agents. All that is required is that each agent expresses a willingness to pay or willingness to accept payment for a stock at a certain price. A decade later, Robert Axtell (2005) demonstrated that decentralized exchange amongst neoclassically rational agents also approximates equilibrium prices and distributions predicted by the neoclassical framework.[13] Axtell finds that there is a linear relationship between the size of a population and the number of interactions, a proxy for computations performed by agents, required to reach equilibrium.

Two features are useful for interpreting agent-based models. First, in a static context, approximation of equilibrium only requires definition of a willingness-to-pay and a willingness-to-accept among agents. We prioritize this requirement in our construction of heuristic decision-making agents (Gigerenzer 2008; Kahnemann 2011). Second, the number of interactions (i.e., exchanges) among agents provide a meaningful measure of time. A similar insight is drawn by those who trade algorithmically observing time in terms of trading volumes (Easley and de Prado 2012).[14] Markets process information. They perform this function through prices that emerge during the course of production and exchange. Information transmitted by prices is integral to the process of equilibration. Evolutionary models with heuristic agents clarify insights from Alchian (1950) and from Hayek's expansive research program that both recognized rationalization of the economic activity to be the result of refinement by competitive forces that influence and are coordinated by the price system.[15] An agent-based model is an ideal tool for demonstrating learning within dynamic economic processes.

To frame search within the market algorithm, we draw on the metaphor of the evolutionary landscape provided by Stuart Kauffman and Simon Levin (1987). Kauffman and Levin developed a model of peptide evolution. Uncorrelated fitness landscapes demonstrate the relationship between learning and fitness improvement. A landscape consists of 2^N vertices, thus representing all possible states of the peptide chain. Solving for optimal fitness on the resultant NK landscape is just a combinatorial problem to be solved by random mutation of peptides from local optimum. The number of local optima on a landscape is defined by K, with K = 0 representing the case of a single global optimum and no other local optima. With discovery of every new, higher local optimum, fewer and fewer improvements remain available. Consequently improvement from one local optimum to the next will take longer to find. Kauffman and Levin show that, for a limited level of complexity, the log value of their fitness criterion improves linearly over the log of generations created. The authors generalize the model to a correlated landscape by building on basic intuition that bears some resemblance to a binary search algorithm.

Kauffman and Levin generalize their model to a correlated landscape and Kauffman and Johnsen (1991) include coevolution of peptide chains with binary values. Valente (2014) provides a method of including continuous variables on a static landscape, but none of these models include coevolution. In all cases, even when the landscape analyzed is relatively large, the number of local maxima and the structure of

[13] In our model, we call these neoclassical agents Optimizers. These agents use utility functions with constant elasticity of substitution.

[14] Although Mandelbrot promoted the idea of trading time, he did not equate this with measuring time in terms of trade volume (Mandelbrot and Hudson 2004).

[15] For an informative perspective concerning Hayek and Complexity, see Axtell (2016)

coevolution is prestated and could be revealed by one who constructs the model. Within a market, agent evolution is interdependent, and the parameter space explored is vast. The fitness of the economic system across this space can only be revealed via computation within the market process (Axtell 2005).

Our genetic algorithm employs dynamic meso-level structures, transforming rules that guide agent decision-making and the parameters these rules employ (Dopfer, Foster, and Potts 2004). As a result, this model does not predefine the structure of the landscape in terms of K, the number of local maxima. Rather, the landscape emerges as a result of population growth and mutation of agent strategies. We can only infer general features of the landscape from the distribution of outcomes generated over numerous simulations. This moves analysis in the direction of Kauffman's later work that is concerned with uncertainty and creativity. The problem of economic production and exchange is not correctly framed as a problem where "all (or even some) solutions are somehow listed, listable, calculable, or comparable" (Felin, Kauffman, Koppl, and Longo 2014, 273). The only way to accurately forecast future states of the model would be to execute the simulation. The vast domain of possible strategies, while already large for a single agent in a fixed landscape, experiences a combinatorial explosion when you include the strategies, positions, and other attributes of other agents. The simple model of peptide chains exploring landscapes whose features are already known to the scientist, while not irrelevant, fails to capture the uncertainty facing the modeler. For example, even with only 5 parameters under consideration, Alkemade et al., (2009, p. 535) find that their proposed space of exploration for development of energy systems includes "2,646 theoretical combinations of energy sources". Adding only a small number of parameters greatly increases this complexity, especially as those parameters take on continuous values. And while selecting a limited set of fixed parameter values to govern agent decisions may be informative, "values of these parameters have to be carefully chosen in order for the EA [evolutionary algorithm] to perform well" and "premature convergence [about an equilibrium] restricts EA in it's [sic] learning capabilities" (Amman, Poutre, and Alkemade 2006). By allowing the values of these parameters to be selected from distributions and to increase or decrease incrementally, usually without strict bounds, we circumvent the problem of brittleness. This increase in model complexity requires that model results must be interpreted abductively as the economic significance of parameters can only be inferred by outcomes (Axtell and Guerrero 2019; Alkemade, et al., 2007, 164-165). Carefully observing the relationship between outcomes and these complex arrangements of parameters, we account for the significance of computation that "is abstracted away in the standard economic model", but that economists modeling artificially intelligent agents "must account for to operationalize rationality" (Parkes and Wellman 2015, 268). As we are concerned with the trajectory of macroeconomic performance facilitated purely by competitive forces, our heuristic agents employ relatively myopic strategies. Even absent a mapping of expected values across the decisions space, these agents generate aggregate results comparable to optimizing agents.

The economy will generate some of the general features observed by Kauffman and Levin (1987), namely that fitness increases linearly in log-time, thus indicating that metaphor of evolutionary landscapes is useful for interpreting the model's results. By allowing for such a vast space of parameter choices and agent reproduction and failure, we allow for persistent learning and growth that does not lead to swift convergence to an equilibrium level of output or population that is typically observed in evolutionary agent-based models (for example, Vriend 1995; 2001). Although agent-based models have focused on macroeconomic topics, our work employs macroeconomic theory as a straightforward interpretive device to measure the efficiency of heuristic (non-optimizing) agents (for example, see Gabardo et al., 2020; Dosi and Roventini 2019; Gualdi et al., 2015). One exception includes Lengnick (2013) who provides a simple market of firms and labour, however, the author does not focus on efficiency improvements as “[p]roduction technology is also fixed” in his model and is represented by a productivity parameter (105). Productive efficiency is an emergent outcome in our model that we use to measure the fitness of economic arrangements.

Model

The market algorithm facilitates learning over a search of a dynamic landscape where choices of strategy for production and exchange influence the profits and wealth of each agent. Hayek (1968) reminds us that competition is a discovery process. Competition within markets yields an intensive search for ways to serve consumers at the lowest cost. We can think of the fitness of a market in light of its ability to create value for consumers. The fitness of a given firm is indicated by its level of profit and the level of wealth to which that profit contributes. The fitness of markets in our model, which includes only capital and no labor, is indicated by its total productivity. This value is influenced by the average productivity of a unit of capital. Greater efficiency yields greater productivity that supports profitable expansion of the stock of productive capital that further augments total productivity.

Portraying the market as a genetic algorithm requires that we define where in the market process 1) mutation and 2) crossover (partial duplication of agents traits) occur. In a traditional genetic algorithm, agents are typically created to optimize some value on a fixed landscape. We allow crossover to occur over particular exchanges, while mutation may occur during firm reproduction. In the marketplace, entrepreneurs learn from one another on-the-fly. And they learn imperfectly. All that is required for the metaphor of crossover in a genetic algorithm is that entrepreneurs emulate the attributes of those entrepreneurs that they believe are best and that some of these attributes mutate. Since learning happens in the course of exchange, agents in this model copy traits of the wealthiest trading partner with whom they have traded at some rate defined by the crossover rate. This rate is exogenously determined in our model.

Finally, to include evolutionary selection – the last feature indicated by Alchian – the model must allow for failure. Entrepreneurs exit the market when they cannot afford to continue operating. This exit tends to remove the worst performing strategies. Unlike traditional genetic algorithms, fitness of a particular agent cannot be straight-forwardly, or even perfectly, evaluated since the landscape explored by agents is dynamic. *Any* change in an attribute of one agent or another may impact the ability of either or both to accumulate wealth and, therefore, to survive. Further, fitness is not only attributed to the individual agent but to the system itself. We approximate fitness by intensity of capital development and wealth accumulation.

We develop these features within a Heuristic rendition of Sugarscape along the lines demonstrated by Caton (2017; 2019). Our model is constructed using Python and includes several distinctions, including allowance for neoclassical utility maximization, use of only two - rather than three - primary strategies, and the use of target ratios of good stocks rather than target levels.[16]

Mutation

[16] For more information concerning to Sugarscape, we refer to the reader to the mentioned articles and Epstein and Axtell (1996).

Most computational literature recommends specifying a universal and small mutation rate, between 0.001 and 0.05 (Lin, Lee and Hong 2003). This is comparable to the small mutation rate observed in biological systems (Sims 1994). When initially defining the market in this model, the search space is unexplored. To allow for sufficient search of all parameters, and to allow for a market that has some agents with higher r, each agent's rate of mutation is drawn from a distribution that includes relatively high rates of mutation. Agents with higher rates of mutation are relatively more creative as a higher rate of mutation promotes greater levels of experimentation.

Agents in the model are assigned a rate of mutation, r , that shapes the extent of genetic variation (De Jong, 2006). The parameter r determines the probability that an agent deviates from its parent's attribute values *and* the magnitude of such deviation. An initial distribution with *too many* agents with extreme r values could result in instability. The solution is to initially distribute mutation rates using a reciprocal uniform distribution. The initial distribution of r values are distributed between lower boundary r_{min} and upper boundary r_{max} as:

$$u \sim U(a, b)$$

$$r = (\frac{1}{u})$$

In this model, $r_{min} = \frac{1}{b}$ and $r_{max} = \frac{1}{a}$ where a = 2 and b = 50 (Figure 1). This bounds initial mutation rates between 0.02 and 0.5. Values drawn from this distribution are close to the lower bound, r_{min}.

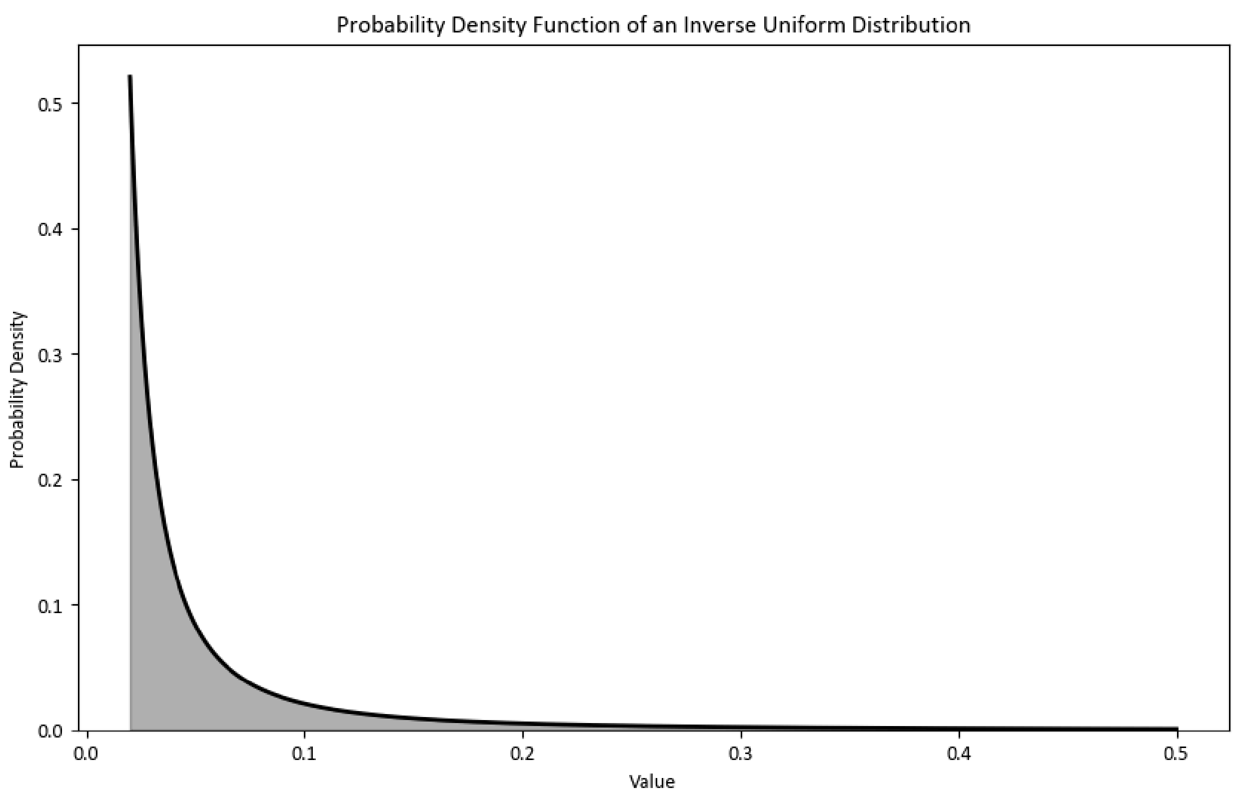

Figure 1

Allowance for high rates of mutation follow Caton (2019) where r governs the probability of a mutation and the severity of the divergences resulting from mutation. For a mutate-able attribute, a, with parental attribute a_p, $E(a) = E(a_p)$. Mutation of parameters whose values exist across a continuous domain mutate at a rate of r (i.e., for each child parameter, if $x \sim U(0, 1) < r$, then mutate). The parameter, a, unmutated, is the value, a_p, inherited from the parent. If the value a is mutated from the initial value, a_p then mutation is described as follows:

1. Determine direction of mutation for a according to a binomial random variable φ, which can take on a value of 1 or -1, each with a probability of 0.5.
2. Mutate value based on direction – changes must be symmetric:

$$a = e^{ln\left(a_p\right) + \varphi ln\left(a_p r\right)} \quad (4)$$

For Boolean variables (i.e., breed and herder strategy), mutation also occurs with a probability of r. Mutation switches a Boolean value from True to False. Thus, an agent's primary strategy switches between Basic and Optimizer. The secondary Herding strategy can be switched on or off.

An agent's attribute, r, is itself mutate-able, allowing for the extent of search to vary, guided by incentives generated within the environment. This draws on the logic of Caton (2019) and ensures exhaustive search space coverage initially, followed by convergence once a streamlined strategy emerges. This mechanism enables the market to regulate the extent of creativity. While such deviations from established norms may result in failures, occasionally, these deviations result in strategies that improve agent performance. Higher quality strategies are adopted by a combination of reproduction and herding behavior.

Strategies

Strategies are divided into two classes. There are primary strategies that include the Basic and the Optimizer strategies. The Basic strategy is composed of heuristic evaluations that guide production and trade. The Optimizer strategy is derived from the neoclassical calculus of optimization, as presented by Epstein and Axtell (1996) and Axtell (2005). The secondary strategy, Herding, can be employed by either kind of agent and serves as a means of information transmission between exchanging agents.

Agent Valuation Strategies

Optimizers

In the original Sugarscape, optimizers hold an internal utility function (called a welfare function in GAS), defined as:

$$U\left(w_1, w_2\right) = w_1^{m_1/m_t} w_2^{m_2/m_t} \quad (5)$$

Where w_1 and w_2 are the accumulations of the two goods and m_1 and m_2 are the metabolism of each good per period. The agent's level of wealth can then be described as "the amount of time until death given no further resource gathering" (GAS p. 97) defined as the level of the good divided by its consumption rate. Wealth of each good is denoted $\tau_1 = \frac{w_1}{m_1}$, $\tau_2 = \frac{w_2}{m_2}$. This utility function is of a Cobb-Douglas (1928) form. Agents do not expend their entire resource bundle each period. The valuation of a marginal good for a neoclassical agent in this model is determined by their Marginal Rate of Substitution (MRS) (GAS p.102):

$$
\begin{aligned}
MRS &= \frac{dw_1}{dw_2} \\
&= \frac{\frac{\partial U}{\partial w_1}}{\frac{\partial U}{\partial w_2}} \\
&= \frac{\frac{m_1}{m_t} w_1^{\frac{m_1 - m_t}{m_t}} w_2^{\frac{m_2}{m_t}}}{\frac{w_1}{m_t} w_1^{\frac{m_1}{m_t}} w_2^{\frac{m_2 - m_t}{m_t}}} \\
&= \frac{m_1 w_2}{m_2 w_1} \qquad (6) \\
&= \frac{\frac{w_2}{m_2}}{\frac{w_1}{m_1}} \\
&= \frac{\tau_2}{\tau_1}
\end{aligned}
$$

This form, along with the simplicity of a model containing only two goods, makes MRS a calculation involving just a few algebraic steps after skipping the intermediary calculations to arrive at the final valuation: the simple wealth ratio of τ_2 and τ_1, water and sugar, respectively. This simplification greatly improves the computational efficiency of optimizer decision-making.

MRS represents the relative scarcity of each good for the agent. An agent with an MRS of greater than 1 believes they are relatively rich in water and poor in sugar, and therefore should target sugar. An agent with an MRS of 1 values sugar and water equally. An agent with an MRS of less than 1 is poorer in water and thus will target water. This allows us to exclude the perfect information assumption that typically describes neoclassical optimization. Agents only consider local information and only make local trades. They are boundedly rational utility maximizers.

Basics

In keeping with a model of agents who simplify valuation decisions, Basic agents do not continuously calculate marginal utilities. They instead hold a reservation ratio, γ. This value represents a simple, genetically determined heuristic describing the ratio at which agents would like to hold water to sugar. That is, in this simple context of two goods, γ takes the same form as MRS – a scalar ratio. Importantly, this formulation would not be possible in a system containing more than two goods. In a market with n goods, Basic agents would hold their heuristically determined marginal valuations of each good as a unit vector of length n.

An agent's initial choice of γ is drawn from a lognormal distribution across the population, centered at 1. It evolves genetically through heritability and mutation. A γ of 1 would imply that the agent would like to hold equal amounts sugar and water, a γ of 2 would imply the agent values water twice as much as sugar, and a γ of 0.5 would imply the agent values sugar twice as much as water. γ can change according the preference change heuristic, δ. This represents the rate at which the agent adjusts γ in light personal excess supplies and demands for each good. The mechanics of δ will be discussed in a later section.

If an agent discovers a γ and δ combination that gives them a reproductive advantage in their geospatial context and surrounding agents, they will reproduce more often, and that strategy will percolate to nearby agents. Optimizers form valuation decisions from their calculated marginal utilities. Basic agents do not recalculate their preferences every time their bundle of goods changes. They adjust δ at the end of each period.

Herding

Basic and Optimizing agents might employ a Herding strategy. Herding agents emulate the strategies of their more prosperous peers, as reflected by fitness metrics, either wealth or reproductive success. Herding is ubiquitous in real-world markets (Banerjee 1992). Selection of an agent's fitness indicator, by which trade partners will be judged, is not obvious. We could let agents experiment by randomly choosing various indicators. However, we assert that agents may be interested in one of two variables. Agents may herd attributes from wealthier trading partners. Both goods are required for reproduction, and an agent with high wealth is not guaranteed to have achieved high wealth using a strategy that will allow for high levels of reproduction. This is because high stocks of a single good can result in high agent wealth, even if the agent will exit the market soon due to a lack of the other good. Alternately, agents can track reproductive success of their trading partners. If an agent identifies a peer with greater prosperity, i.e., more alive children (more successful expansion into the market), they adopt the peer's traits. If they see the offspring of a partner that they once thought to be more successful than them fail in the market, they may pivot to herd from a trading partner judged to be more successful.

Herding agents choose from either strategy. An agent's herding strategy can be changed through mutation or herding. Herding agents track the fittest (defined by their herding strategy) trading partner they have encountered, and only herd from agents who are still alive – if their fittest agent is forced from the model, they become their own fittest agent until they encounter a trading partner with a higher fitness rating.

In a traditional genetic algorithm, attributes from the fittest agents are combined after each period through a crossover process. This process gives the child an equal probability of inheriting either parent's trait – they inherit only one of the parents' traits – not both. This process results in an agent with a unique combination of traits drawn only from the selected pool of top performing agents. Our model does not have sexual reproduction. Instead, the model uses an exogenously determined crossover rate, k, representing the probability that a herding agent will copy a given trait of their partner. For simplicity, this value is set at 0.5. On average within a given exchange, half of the traits of a trading partner perceived to be superior will be copied.

Agent Reproduction

Our model diverges significantly from the original Sugarscape regarding reproduction. Instead of modeling agents like biological entities that reproduce sexually, as Epstein and Axtell did, we liken agents to entrepreneurial firms (Foss and Klein 2012) that propagate asexually. In Epstein and Axtell's model, agents needed a partner to reproduce, drawing closer comparisons to biological entities rather than market-based firms.[17]

Basics use a heritable, mutable parameter to set their reproductive standards. This parameter, represented as θ_i, is the reproduction ratio that denotes the multiplier of an agent's initial good endowment they must gather before they can reproduce. Initially, this ratio can vary between 5/4 (where an agent needs only one-fifth of their initial endowment in order to reproduce) to five (where they need to acquire five times their initial endowment to give half of that to their offspring):

$$ln(\theta_i) \sim U(ln\left(\frac{5}{4}\right), ln\,(5))$$

$$where\ U(a, b)\ is\ the\ uniform\ distribution\ from\ point\ a\ to\ b$$

Thus, an agent's complete reproduction criteria in terms of goods are represented by:

$$R_i = \theta_i \varepsilon_i \qquad (1)$$

Where ε_i is the agent's primary endowment of good i. Through the genetic algorithm, basic GA agents are predicted to settle on a reproductive standard that provides a kind of inherent, evolving fitness function. While the form of this reproductive standard remains consistently tied to the product of θ_i and ε_i, it is the hereditary nature of θ_i that enables agents to adjust their reproductive criteria over time. The complete reproduction rule for basic agents in each cycle is as follows:

1. Examine stocks of each good.
2. If $w_s > R_s$ (sugar) and $w_w > R_w$ (water), then reproduce and give offspring half of the reproduction criteria for each good.

Early in a market's development, firms, uncertain of which strategies are most effective, can choose different savings levels before launching new entities. Some might retain too few resources,

---

[17] Their sexual reproduction rule (Epstein & Axtell 1996) involved:

1. Randomly select a neighboring agent.
2. If the neighbor was fertile and of the opposite sex and there is an empty patch on a neighboring site for the offspring, then reproduce.
3. Mendelian rules for genotype selection: for each attribute (metabolism, vision, etc.), choose one of 2n2 equally likely genotypes for the child, where n is the number of attributes parents can pass to children.
4. Repeat for all neighbors.

Typically, in a standard genetic algorithm, agents' reproductive potential is assessed using an external fitness function. In the Sugarscape model, however, the capacity to reproduce is gauged by an agent's ability to accumulate enough goods to pass on to their offspring. The genetic approach blends these concepts, employing a genetically determined fitness function anticipated to hone in on efficient values as the model progresses.

overextending themselves and ultimately failing. Others might delay expansion, getting outpaced by more aggressive competitors.

Basic agents reproduction follows a parameter that indicates the proportion of their initial endowment required for reproduction, θ_i. This ratio is effective in regulating both basic and optimizing agent's ability to find efficient wealth levels for reproduction, eliminating population growth oscillations seen in GAS (p.64-67). Given this, and the primary focus on comparing valuation representations rather than reproductive dynamics, optimizing agents were also endowed with a genetic, alterable reproduction ratio parameter identical to that of the Basic agents, θ_i. *This is the lone trait optimizing agents can pass down to their offspring*.

Agent Movement

Optimizers

In the original Sugarscape, agents chose to move to the patch which increased their utility the most, according to the following function which maximizes their prospective utility, given the good types and quantities on their neighboring patches:

$$U\left(w_1 + x_1^s,\ w_2 + x_2^s\right) \qquad (7)$$

Where s is a site contained within the agent's neighborhood of vision, N_v, and x_1^s and x_2^s are the levels of water and sugar at s.

Basics

The movement rule the basic agents use is a similar utility maximization as the optimizers, but is instead based on their subjective marginal valuation, γ:

$$(\frac{1}{\gamma}x_1^s,\ \gamma x_2^s) \qquad (8)$$

This form allows for the valuation of each site to be based on the agent's relative valuation of water and sugar. If $\gamma > 1$, then the agent prefers water, and sites with water will be given more weight in the calculation. Inverting γ for sites with sugar allows them to be weighted heavier for agents who prefer sugar, i.e., $\gamma < 1$. This significantly improves the Basic agent's ability to manage one's stock of sugar and water.

Basic Preference Adjustment

One of the salient features of basic agents in the GA model is the potential for modulating γ through another variable, δ. In GAS, optimizing agents maintained a uniform preference-adjustment mechanism.

Even though their preferences could adapt in line with fluctuations in metabolism and respective good's wealth, as illustrated by Equation (1), utility functions are homogeneous (Epstein and Axtell, 1996). By contrast, the basic agents not only exhibit heterogeneity in their valuations but also methodologies for adjusting those valuations.

The parameter δ spans a log-uniform range from 0.5 to 2 in the first generation of agents, ensuring that the expected change caused by this parameter is multiplication of γ by 1. In subsequent generations, δ is set in light of inheritance, mutation, and herding. The parameter δ governs the intensity and direction with which an agent recalibrates their reservation ratio, γ toward their ratio needed for reproduction, R_r. To ensure the organic dissemination of this heuristic across geographical landscapes akin to γ,we follow a rule similar to the adjustment mechanism described by Caton (2017). Agents base their reservation ratio adjustment On two factors: 1. The discrepancy between their current goods ratio and their ideal goods ratio (which is defined in equation 9), and 2. Their δ, which governs the intensity of their adjustments of γ (the adjustment mechanism is shown in equation 10). In Caton (2017), the rate of change in preferences is defined as a marginal increment. This model differs in that the adjustment parameter is a proportion of the full adjustment (the full adjustment being the one that would change the agents reservation ratio to exactly match their reproduction ratio). The full preference adjustment process is described next.

Agent Period-Level Reservation Check

Each period, basic agents perform a reservation check, which tunes their preference ratio γ by analyzing the disparity between the ratio of goods needed for reproduction, R_r, and the current ratio of goods ($\frac{w}{s}$), combined with the agent's innate propensity to change their preferences, δ:

$$R_r = \frac{R_w}{R_s} \qquad (9)$$

R_w and R_s are the reproduction criteria for water and sugar, respectively, shown in equation (1):

$$R_i = \theta_i \epsilon_i$$

This leads to the change in γ, which is added to the agents previous γ:

$$\Delta\gamma = \delta(ln\left(\frac{\frac{w}{s}}{R_r}\right)) \qquad (10)$$

These adjustments occur each period.

While the proposed adjustment mechanism is far from being the sole conceivable rule, it is paramount to understand its contextual suitability. While the particular rule selected will certainly shape outcomes generated by the model, any rules will be subject to selection. Competition and selection determine which rules are adopted within a population.

Exchange

Modeling of decentralized trade, without a central clearinghouse or auctioneer, is reminiscent of the search models in monetary economics where trading is frictional and not instantaneous (Lagos and Wright 2005). The model described in this paper continues this decentralized trend but evolves the valuation paradigm from a strictly neoclassical framework.

Optimizers

When two agents meet, they fully reveal their internal valuations (MRSs) to each other. If they have different MRSs, then they can trade. The agent with a higher MRS is relatively richer in water and will trade for sugar. The agent with a lower MRS is relatively richer in sugar and will trade for water. Once the direction of trade is established, agents must agree on a price. The price is a ratio of the two goods, and is calculated using the geometric mean of the agents MRSs:

$$p = \sqrt{MRS_A MRS_B}$$

Table 1: Relative MRSs and the Direction of Resource Exchange for Optimizing Agents in GAS

	MRS(A) > MRS(B)		MRS(A) < MRS(B)	
Action	A	B	A	B
Buys	sugar	water	water	sugar
Sells	water	sugar	sugar	water

*Borrowed from GAS p. 103 (Epstein and Axtell, 1996)

The trade goes forward only if both agents are made better off according to their respective internal valuations. The trade algorithm between two agents in GAS (p. 105) is as follows:

1. Agent and neighbor compute MRSs. If they are equal, do not trade, otherwise proceed.
2. Water will flow from the agent with the higher MRS to the agent with the lower MRS and vice versa.
3. Compute price, $p = \sqrt{MRS_A MRS_B}$.
4. If $p > 1$, then p units of water are exchanged for 1 unit of sugar; if $p < 1$, then $1/p$ units of sugar for 1 unit of water.
5. If each agent can afford the trade, it will make each agent better off, and not cause the agents MRSs to cross over, execute the trade and return to step 1.

Basics

In the GA model, γ replaces MRS as an indicator for direction of trade but works in an inverse fashion due to the form of γ. Additionally, basic agents do not use the preferences of their trading partner as a signal to stop trading, but rather stop when their actual ratio of goods ($\frac{w}{s}$) would crossover their desired ratio of goods, γ. Table 2 shows how the basic GA agents use their arbitrary valuation heuristic, γ, to decide trade direction.

Table 2: Relative γ and the Direction of Resource Exchange for GA Basic Agents

	$\gamma_A < \left(\frac{w}{s}\right)_A$ & $\gamma_B > \left(\frac{w}{s}\right)_B$		$\gamma_A > \left(\frac{w}{s}\right)_A$ & $\gamma_B < \left(\frac{w}{s}\right)_B$	
Action	A	B	A	B
Buys	sugar	water	water	sugar
Sells	water	sugar	sugar	water

This ensures that the price is beneficial for both agents, given their heuristic marginal valuation of each resource. The price for trades between basic GA agents is analogously:

$$p = \sqrt{\gamma_A \gamma_B}$$

Their reservation ratio of Basic agents represents their "ideal" ratio of the two goods. To execute a trade then, they should trade until one more trade would cause one of the two agents (or both) goods ratios to cross over their own reservation ratio:

1. Agent and neighbor compare γ. If they are equal, do not trade, otherwise proceed.
2. Water will flow from the agent with the lower γ to the agent with the higher γ and vice versa – inverse of the mechanism used by optimizers.
3. Compute price, $p = \sqrt{\gamma_A \gamma_B}$.
4. If $p > 1$, then p units of water are exchanged for 1 unit of sugar; if $p < 1$, then $1/p$ units of sugar for 1 unit of water.
5. If each agent can afford the trade and the trade will not cause either agent's ($\frac{w}{s}$) to cross over their own γ, execute the trade and return to step 1.

This trade algorithm reveals how basic agents differ from optimizers. During each trade, optimizers re-compute their MRS to move along the Edgeworth contract curve. For Basic agents, γ does not change during a trade, only at the end of each period through δ. Instead of optimizing according to a CES utility functions, Basic agents marginally adjust γ and evaluate the efficacy of these adjustments based on outcomes.

Exchange Between Basic and Optimizer Agents

Basic agents and optimizing agents can trade with one another. The direction of the exchange of goods in a basic-optimizer trade is described in table 3.

Table 3: Relative γ, MRS and the Direction of Resource Exchange Between GA Basic Agent(A) and Optimizing Agent(B)

	$\gamma_A > \left(\frac{w}{s}\right)_A > MRS(B)$		$\gamma_A < \left(\frac{w}{s}\right)_A < MRS(B)$	
Action	A	B	A	B

Buys	sugar	water	water	sugar
Sells	water	sugar	sugar	water

The price would compute to be:

$$p = \sqrt{\gamma_A MRS_B}$$

Basic agents will continue to trade until one additional trade would cause their goods ratio to cross over their γ. Optimizing agents will continue the trade until the price of the next trade crosses over their MRS.

Summary of Initial Parameter Selection and Model Details

We compare Epstein and Axtell's original Sugarscape, the GA model, and the combination of these. We will refer to these models by the types of strategies that they include: optimizer, basic, and basic-optimizer. For all models, except in the section examining asymmetries in (still fixed) consumption rates, the consumption rates of both water and sugar were set at 1.5 (a value which allowed the model to reach equilibrium in a reasonable amount of time). The vision of each agent was also fixed at 1. Both these static simplifies analysis of the evolution of trading strategies by ignoring other fitness-increasing attributes like vision.

The range of initial endowments given to the first generation of agents is randomly distributed from 10 to 25 for both water and sugar. The reservation ratios of the first generation of basic agents are drawn from a log uniform distribution ranging from 0.5 (valuing water twice as much as sugar) to 2 (valuing sugar twice as much as water). The reproduction ratio for each good was drawn randomly from a distribution of $\frac{5}{4}$ to 5. Each model's minimum mutation rate was 0.02, and maximum mutation rate was 0.5. 0.5 was the crossover rate for herders in each model. Table 4 shows the attributes held by basic agents which can be mutated and inherited, along with their initial distributions and descriptions. In subsequent generations, each of these parameters can adjust upward or downward in increments influenced by an agent's mutation rate.

Table 4: Mutate-able Basic Agent Attributes and Descriptions

Attribute	Symbol	Initial Distribution	Description
Reproduction Ratio	θ_i	$ln(\theta_i) \sim U(ln\left(\frac{5}{4}\right), ln\,(5)\,)$	Multiplier of agent's initial endowment of good *i* needed to reproduce
Reservation Ratio	γ	$ln(\gamma) \sim U(ln\,(0.5)\,, ln\,(2)\,)$	Agents preferred ratio of water to sugar
Mutation Rate	r	$r = \left(\frac{1}{u}\right) where\ u \sim U(2, 50)$	Governs probability and intensity of mutation
Preference Change	δ	$ln(\delta) \sim U(ln(0.5), ln(2))$	Determines how much an agent shifts reservation ratio each period

Note: reproduction ratio, reservation ratio, and preference change are unlogged.

Results and Discussion

Macroeconomic models typically focus on real income per capita, which represents total productivity with a single period, rather than wealth. Our model lacks labor, so we will focus on the level of capital and its productivity. Since our model employs robots rather than labor, our corollary measures are capital, income per capital, and wealth per capital. For a given level of wealth, a higher level of capital results in a lower wealth per capital. But higher wealth per capital enables capital growth. Our analysis focuses on these measures as analogues of fitness criteria. Unlike traditional genetic algorithms, the global level variables are not targeted directly. Instead, they emerge through microeconomic processes intermediated by rules and strategies at the meso-level. Our reasoning is described in further detail in the section title “Total Factor Productivity and Wealth Per Capital”. Wealth represents the stock of resources available for investment.

Measuring Time[18]

When examining the system’s efficiency, it is common to observe changes in welfare measures with regard to “natural time,” denoted by the number of periods. Benoit Mandelbrot argued that the flow of time in financial markets is subject to regimes of higher and lower volatility. That is, times where trading occurs with greater and with lesser rapidity (Mandelbrot & Hudson, 2004). Similarly, “transactions time” offers a nuanced view by unmasking volatility clusters otherwise obscured in natural time. Thus, we also analyze results as measured by exchanges and computation time (Easley, D. and López de Prado 2012).

The markets in the Basic and Basic-Optimizer models emerge from a complex of agent decisions. To compare strategies where computations performed for each exchange are significantly different, computational runtime appears to be a measure of time capable of normalizing time between models. Decision-making within a single period is more complex for an Optimizing agent than for a Basic agent. A period, thus, requires a greater level of computation for an Optimizing agent than for a Basic agent. Computational time thus accounts for the time costs of decision-making and action that are not captured by using natural time.

We compare results from all models. For each model, we present results from 100 simulations. Simulations that include Optimizers are run for 60,000 periods. This includes the mixed run with Optimizers *and Basics*. Simulations that only include Basics (no Optimizers) occur across 100,000 periods. This selection of simulation length roughly equates the total runtime for each model – an average of 3,941 seconds per run for Basics and 3,907 seconds per run for optimizers. This also roughly equates the number of total exchanges. Conveniently, where performance is measured according to runtime fitness measures appear to converge. For each analysis, the aggregate (mean) model measure at each junction, along with a “cloud” containing each of the 100 individual runs, is plotted against natural time (periods) and cumulative computational runtime.

[18] Further discussion of *transactions time* and *computation time* are provided in the appendix.

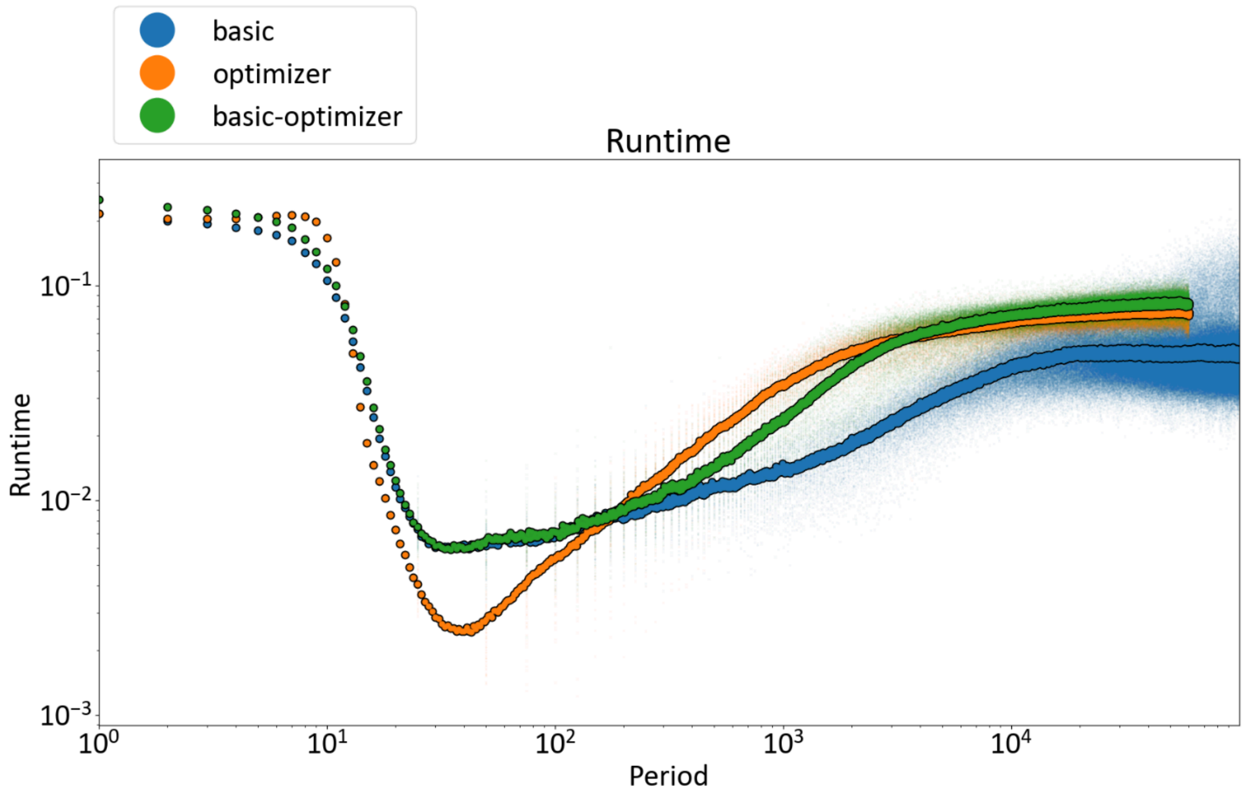

Figure 2: Average Per-agent Runtime by Period, by Breed Set

Capital Levels

The concept of carrying capacity (population) is central to understanding a market's robustness as "a given environment will not support an indefinite population of agents" (Epstein & Axtell 1996, p.30). Since a sustained growth in the level of capital requires improvement in productivity, sustained capital growth indicates improvement in the stock of collective knowledge. As we discuss later, this improvement is not considered by the canonical Solow-Swan model (Solow 1956; 1957; Swan 1956). These improvements are meaningful as they depend on productivity gains in order to be sustained.

We populate all models initially with 2000 agents (out of 2500 available patches), ensuring an expansive exploration of the strategic domain. Nonetheless, due to adversities like unfavorable geospatial conditions or suboptimal strategies, a significant number of agents perish within the first 50 periods. Survival selects for agents that employ relatively more efficient strategies.

Strategic diversity amongst Basic agents prevents capital levels from reaching significantly lower levels during this initial learning period. Those Optimizing agents that do survive, however, quickly outperform Basic agents. On a log scale, the level of capital in simulations with only Basics approximately converges with the level of capital in runs with optimizers. Capital levels in Basic only simulations are more volatile.

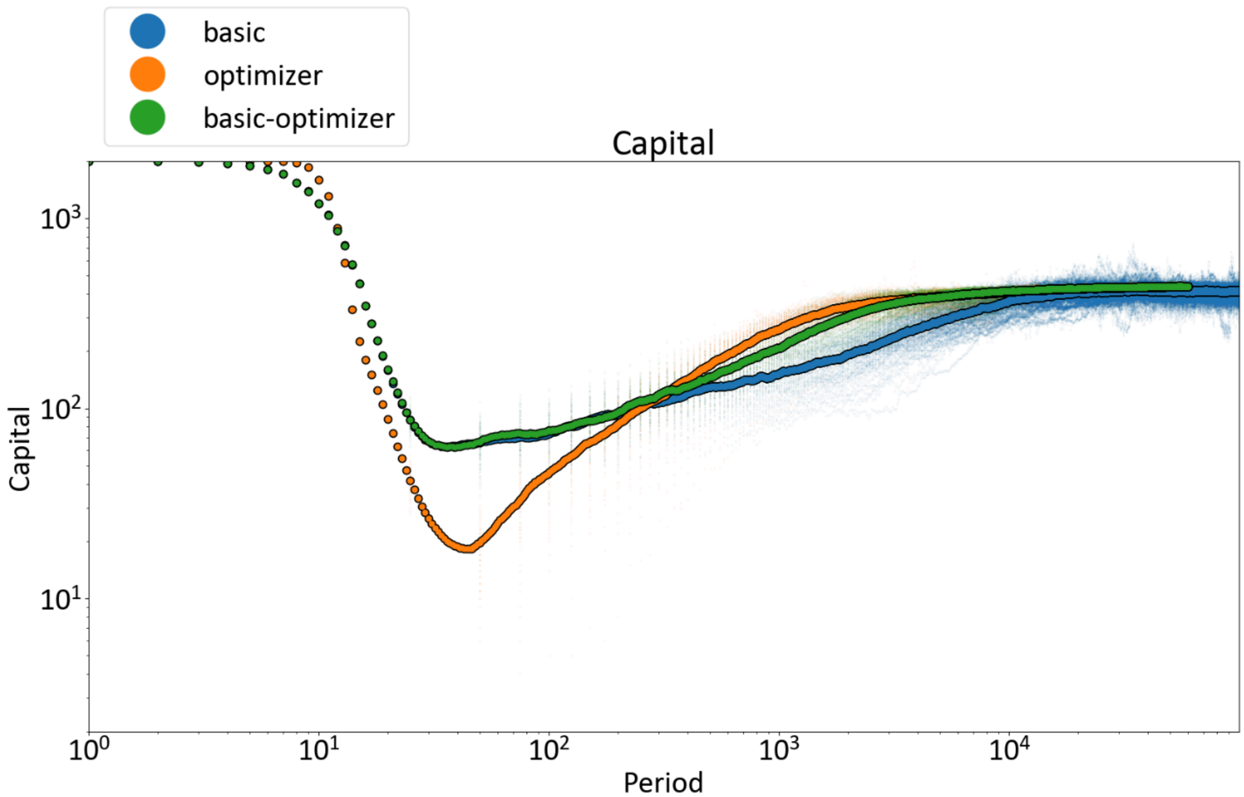

Figure 3: Logarithm of Capital by Logarithm of Period, by Breed Set

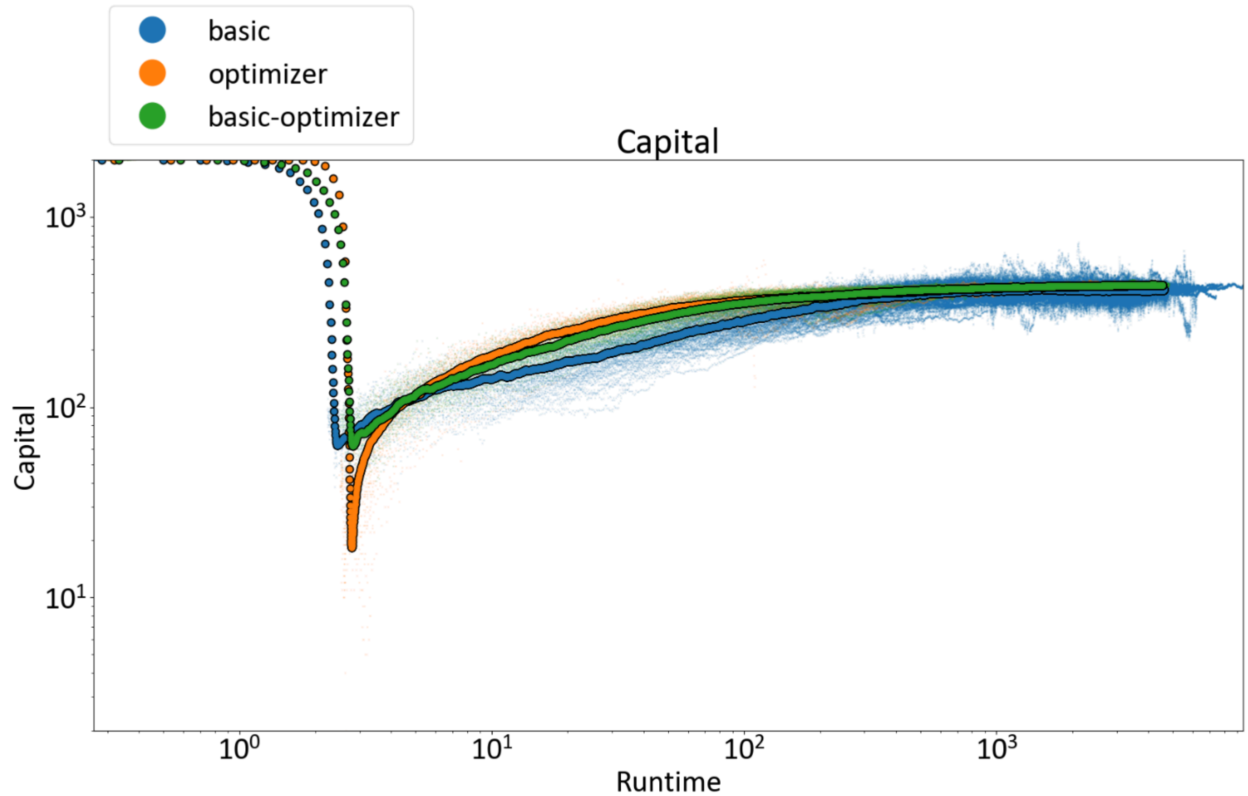

Figure 4: Logarithm of Capital by Logarithm of Cumulative Runtime by Breed Set

Total Factor Productivity and Wealth per Capital

We model total factor productivity to understand improvements in efficiency of the capital stock. In modern macroeconomics, the canonical aggregate production function takes the form:

$$Y = AK^{\alpha}L^{\beta}$$

Where Y is aggregate production per period, K is capital input with contribution α, L is labor input with contribution β, and A is total factor productivity. A is a measure of the efficiency of inputs. A higher A would signal a more efficient economy which could produce more than an economy with a lower A, given identical inputs.

In the Sugarscape economy, there is only one input, capital. Decision-making agents can be thought of as firm owners in control over the wealth in their possession. The value of the firm is equivalent to the resources under their control. We therefore treat each firm as a unit of capital whose existence depends upon the sustainable extraction of resources. Capital receives 100% of returns production factors:

$$Y = AK$$

We thus define total factor productivity, A, as:

$$A = \frac{Y}{K}$$

The Sugarscape economy does not begin in equilibrium. Thus, the implied measure of A is out of equilibrium. As an economy approaches a new steady state, the marginal profit from investment reaches zero. In order for the capital stock to grow, the average productivity of the capital stock must increase sufficiently to promote such an increase. Notice that this is modestly different than the textbook presentation of the Solow model where A is exogenously determined and its estimate does not change as the capital stock grows. With labor in the model, growth of real income per labor grows with the effective capital stock, $\frac{Y}{L}$. However, without labor, improvement in capital efficiency is simply reflected in capital growth. More productive capital generates net improvements in wealth that enable a sustainable increase in the capital stock. In a frictionless equilibrium, income per capita would be exactly sufficient to meet agent consumption requirements.

When capital is less efficient, a greater value of $\frac{Y}{K}$ is required to sustain a given level of K. As capital becomes more efficient, we should expect $\frac{Y}{K}$ to fall as it converges with that rate of wealth consumption. As A converges to this long run equilibrium, improvement in efficiency is solely represented by a growing capital stock. Increases in income are saved until those savings are sufficient to generate a new unit of capital. For the value $\frac{Y}{K}$, the creation of a new unit of capital hides this improvement in average productivity as the consequence of greater efficiency is an increase in the value of the denominator.

Given the simplicity of the model and a lack of monetary medium, we are left squarely in a world described by Say's Law. Demand is supported by an increase in supply. Consumption in excess of production leads to a net decrease in wealth as there exists no monetary disequilibrium where excess demand for money generates an excess supply of goods (Clower and Leijonhufvud 1973). Measures of real income sum consumption and savings. However, savings is treated as always being positive or a value of 0. Contractions in the stock of wealth are treated as consumption and thus positively contribute to real income. Income is thus a flow that does not consistently correspond to changes in the stock of wealth. Income per capita must, on average, exceed the per period consumption requirement. When income per capita exceeds the consumption requirement, population grows, thus leading income per capita to fall. If it is less than the per period consumption requirements, population will fall until income per capita at least matches this requirement (Figure 9). Likewise, the lower bound of savings converges to zero. There does exist some variance, however, as agents do not have perfect information or costless access to trading partners, as compared to Caton (2017). This is especially true for Basic agents who appear to need more savings to support a given level of capital (Figure 8). At the highest levels of capital observed in the simulations, the level of savings amongst basic agents is about an order of magnitude higher than among Optimizing agents. In all models there is a persistent increase in the level of wealth per capital from early in the model. This might be surprising as an increase in capital increases the denominator. If wealth per capital increases as capital increases, this means that the level of wealth increases at a faster rate than capital is increasing (Figure 5). There are two ways to interpret this increase: 1) an increase in wealth per capital indicates an increase in productive efficiency or 2) there could be increases in wealth per capital that indicate potential to improve productive efficiency that would reduce the stock of wealth required to sustainably support each agent. Case 1) appears to be indicated by the general increase in wealth per capital across all 3 runs. Productivity efficiency improves as capital increases. Yet case 2), which seems particularly relevant for simulations that include Basic agents, allows for a decrease in wealth per capital late in the run as capital levels nearly converge to the path from simulations with Optimizers (Figures 6 and 7). Innovations in production and exchange enable the discovery and exploitation of profit opportunities. As the capital stock grows, these opportunities are increasingly scarce, thus requiring greater and greater increases in the wealth stock to support a marginal increase in the stock of capital. This is consistent with the increase in computational costs required for an increase in fitness on fixed evolutionary landscapes.

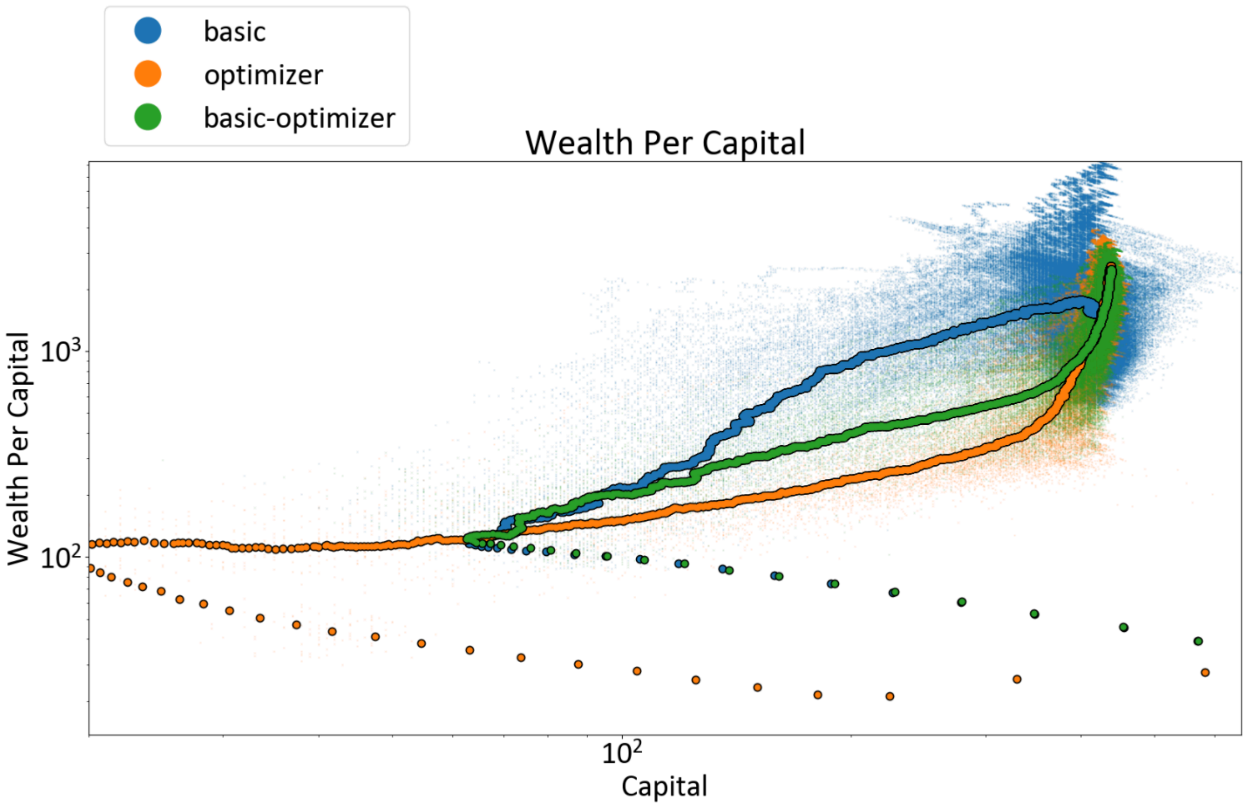

Figure 5: Logarithm of Wealth Per Capital by Logarithm of Capital

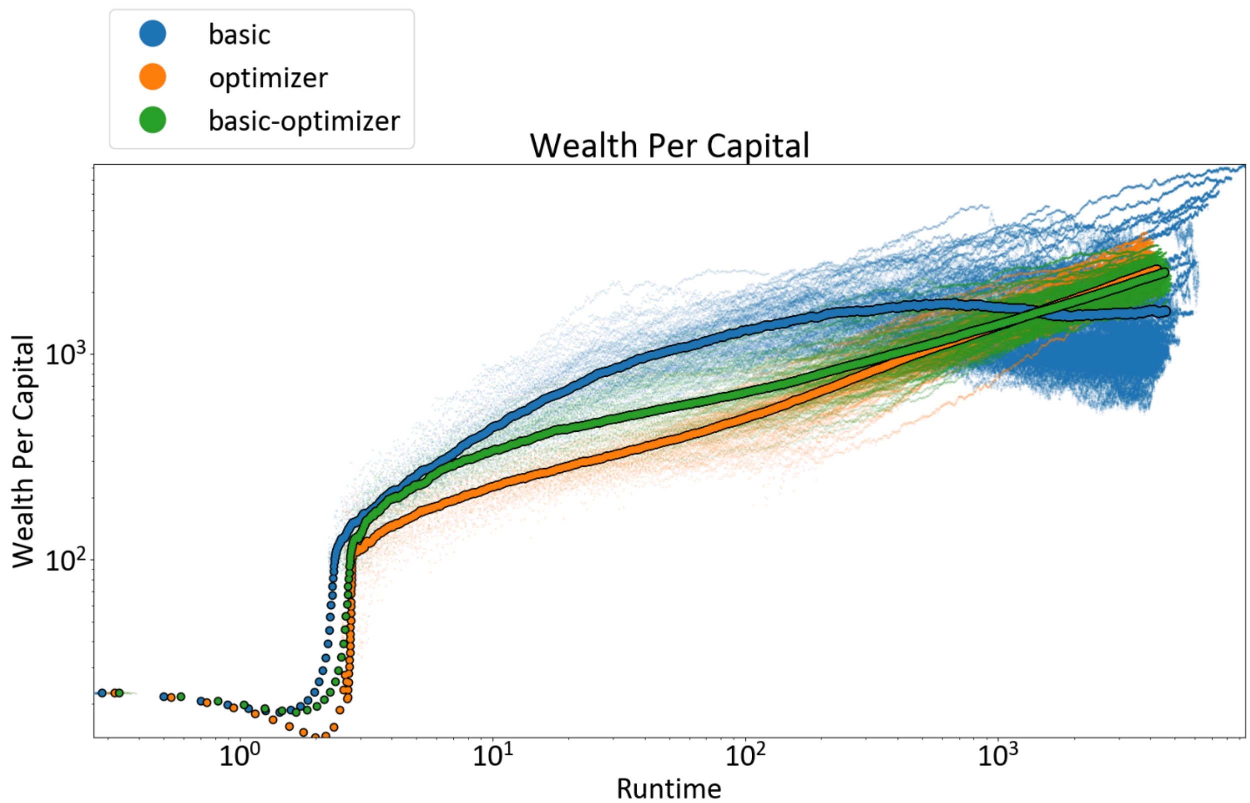

Figure 6: Logarithm of Wealth Per Capital by Logarithm of Runtime

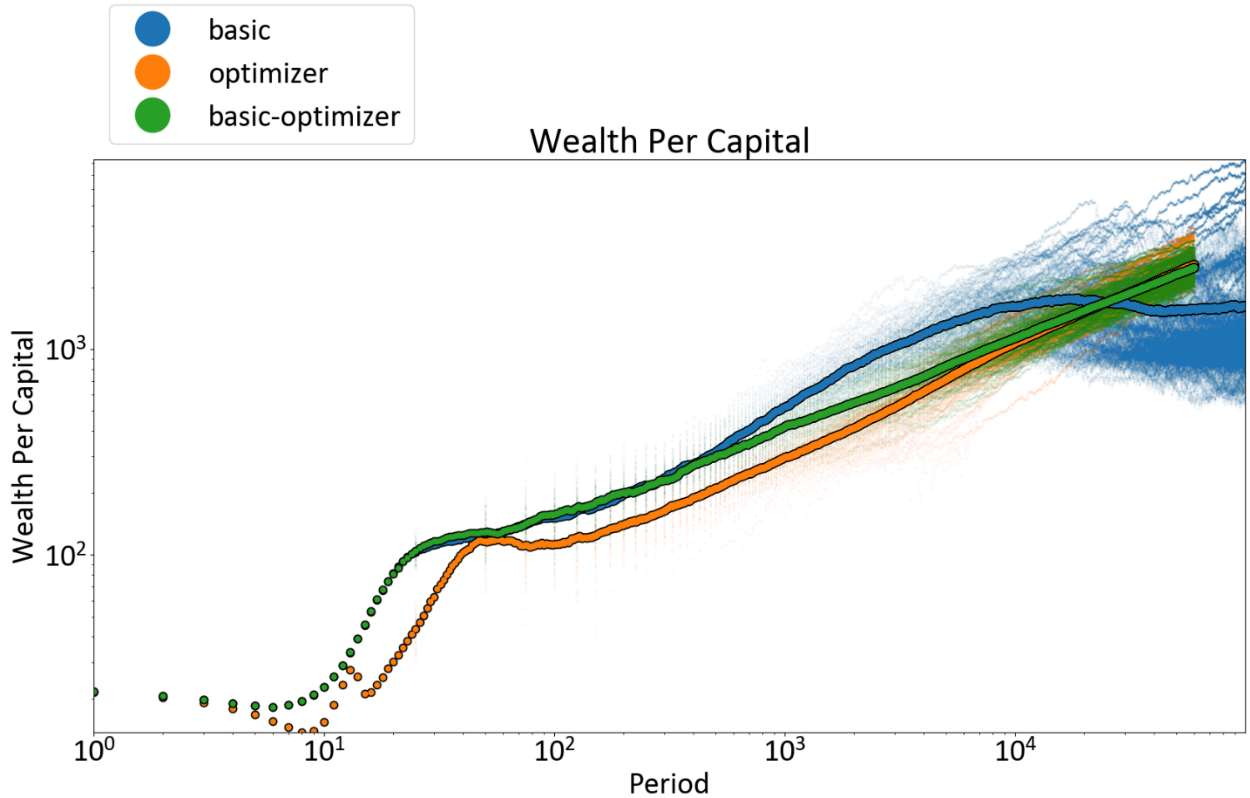

Figure 7: Logarithm of Wealth Per Capital by Logarithm of Period

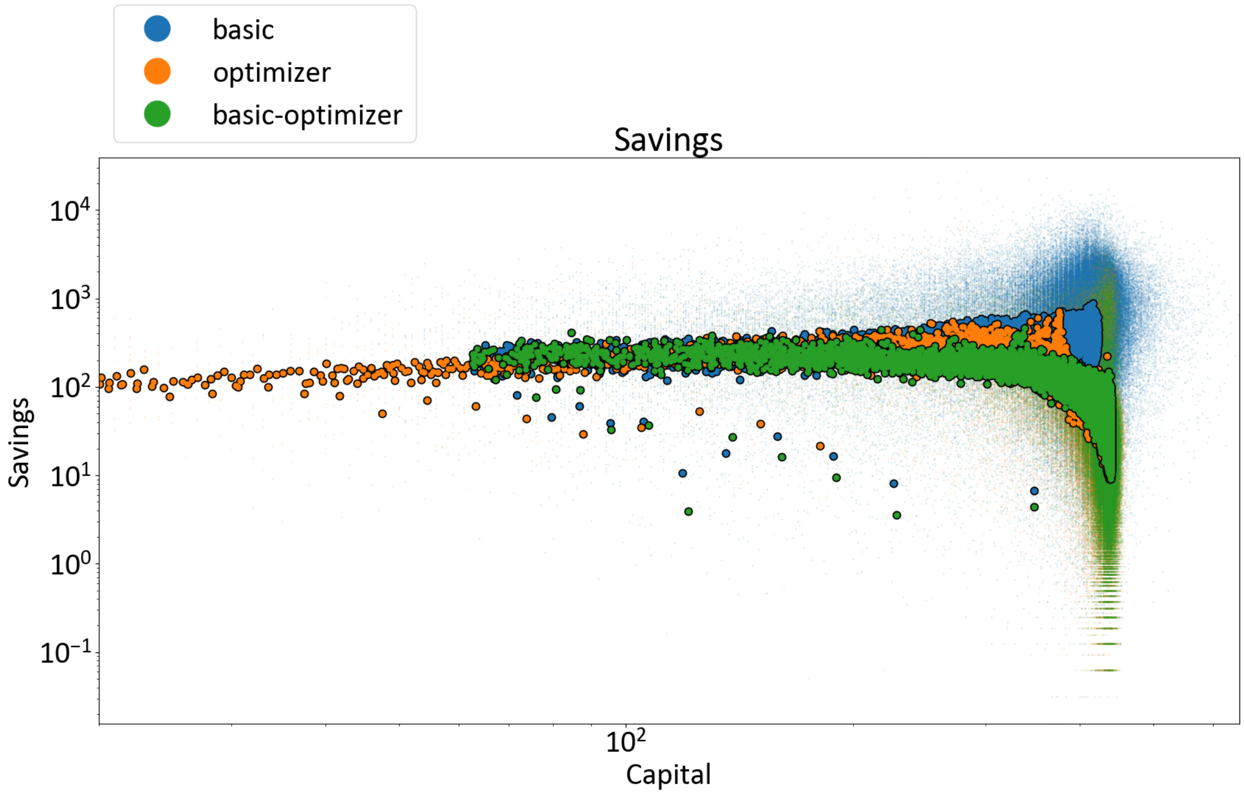

Figure 8: Logarithm of Savings by Logarithm of Capital

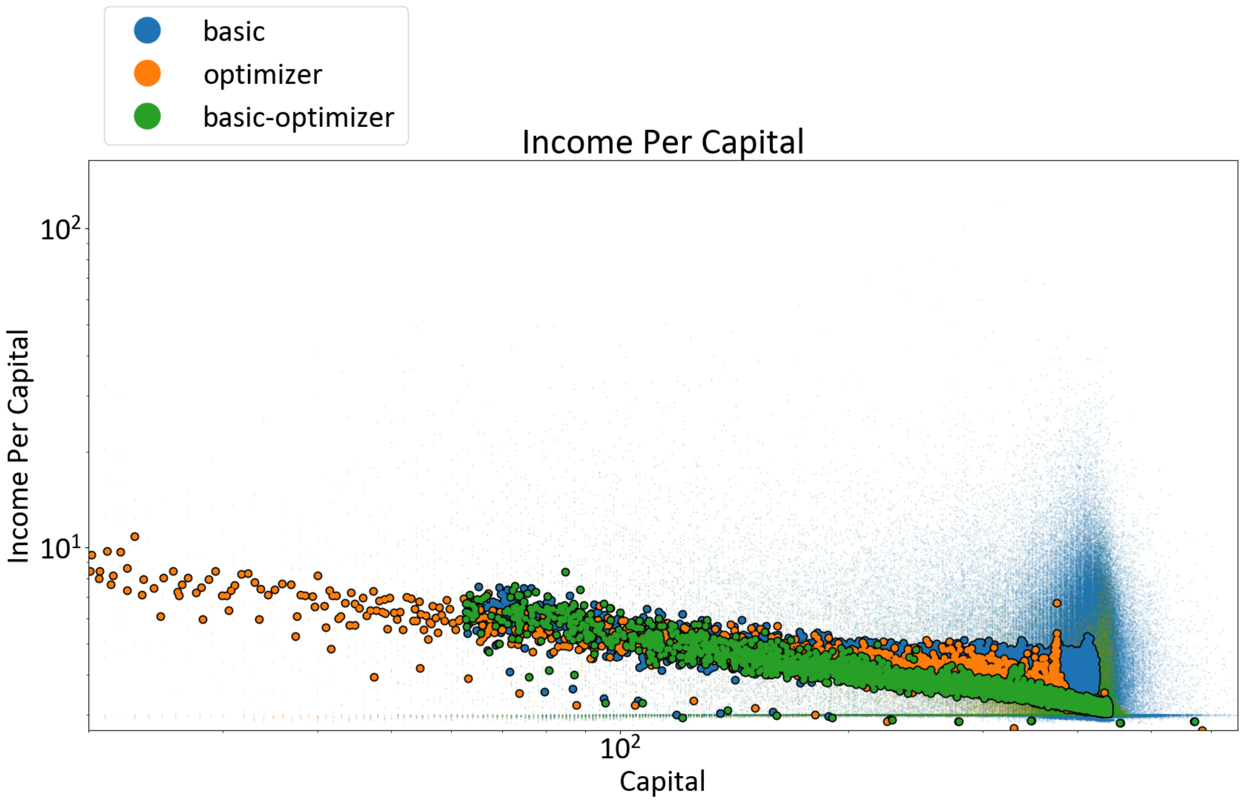

Figure 9: Logarithm of Income Per Capital by Logarithm of Capital

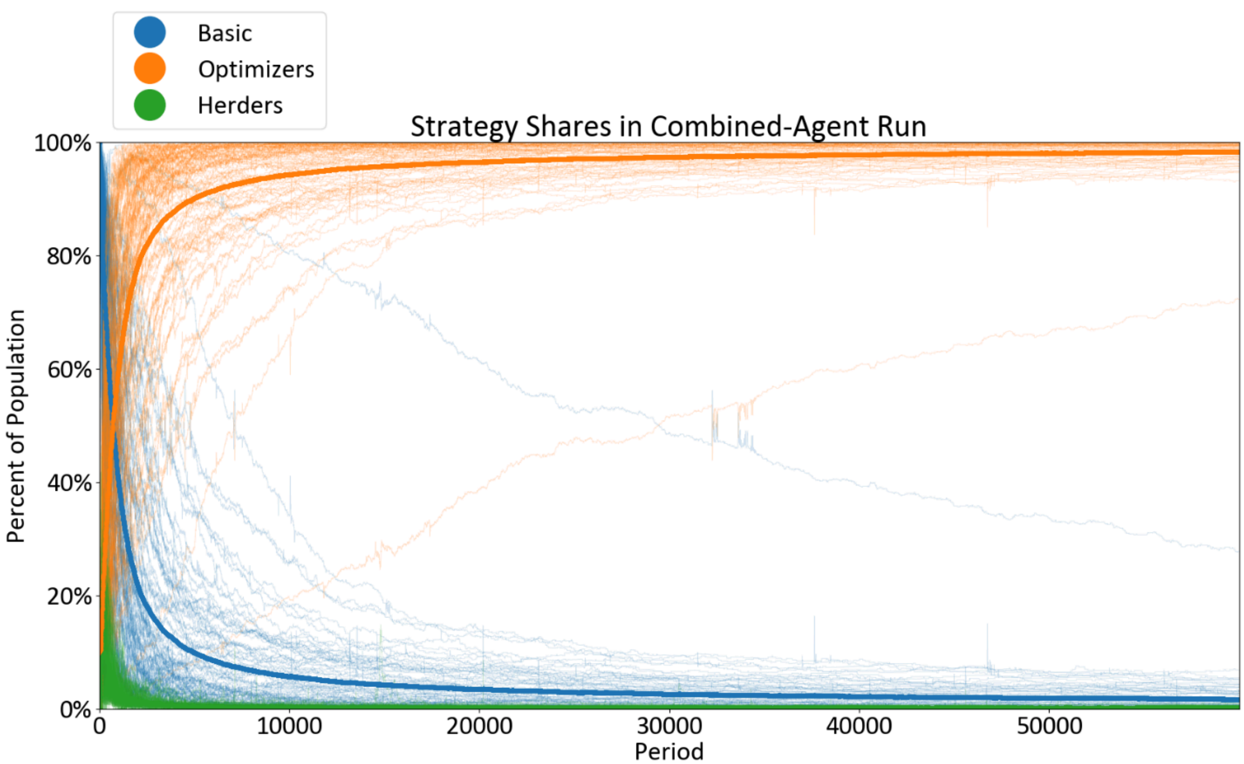

Figure 10: Population Share of Optimizers, Basics, and Herders in Combined Runs Containing Basics and Optimizers

Herder Performance

In each of the simulations, agents are allowed to mimic the attributes of their better performing trading partners, thus allowing for the spread of successful strategies. The only trait that Optimizing agents can herd (or inherit) is reproduction ratio. This, in theory, would allow them to converge to an optimal value for reproduction ratio quickly. This may be why it is observed that runs containing exclusively or mostly Optimizers have much fewer herders (Figure 10). Once the population converges upon a reproduction ratio, the value of searching the parameter space is limited, as is the value of herding. Runs with Basic agents utilize herding to learn efficient strategies since there are many parameters that may be selected, so the set of optimal strategies is continually in flux. The optimal parameter space for a given state of the model is significantly larger for Basics. And the optimal arrangement changes as the number and mixture of parameter settings change.

Basic agents herd more than Optimizers. However, their rate of herding still declines over time as the efficient level of capital approaches its maximum. However, this complacency may lead to stagnation as we see high levels of volatility in the level of wealth per capital as the level of capital in Basic simulations approaches the level of capital in Optimizer simulations (Figure 5-7). Initially, the proportion of herders spikes as it is beneficial to copy more successful partners when it is not obvious which strategies are most beneficial (Figure 11). As the market discovery process generates quality knowledge sets, the value of search declines as do herding and mutation rates (Figures 11 and 12). Finally, Herding agents may follow one of two herding strategies: selection by the maximum number of an agent's living offspring or the maximum level of wealth. In the Basic-only runs the child counting herder strategy seems to outperform the wealth tracking strategy (Figure 13). Both strategies are present throughout, but child counting is the more popular choice for herders in each period.

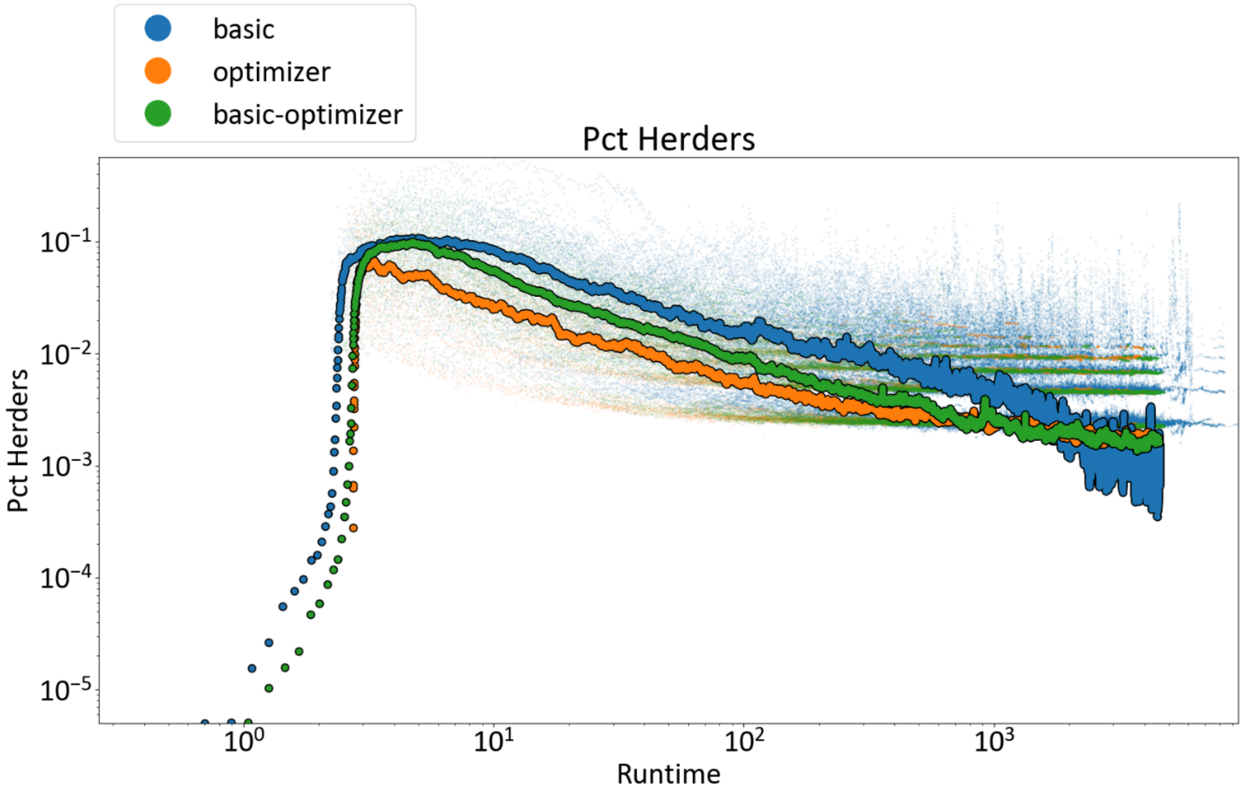

Figure 11: Herder Population Proportion by Logarithm of Cumulative Runtime, by Breed Set

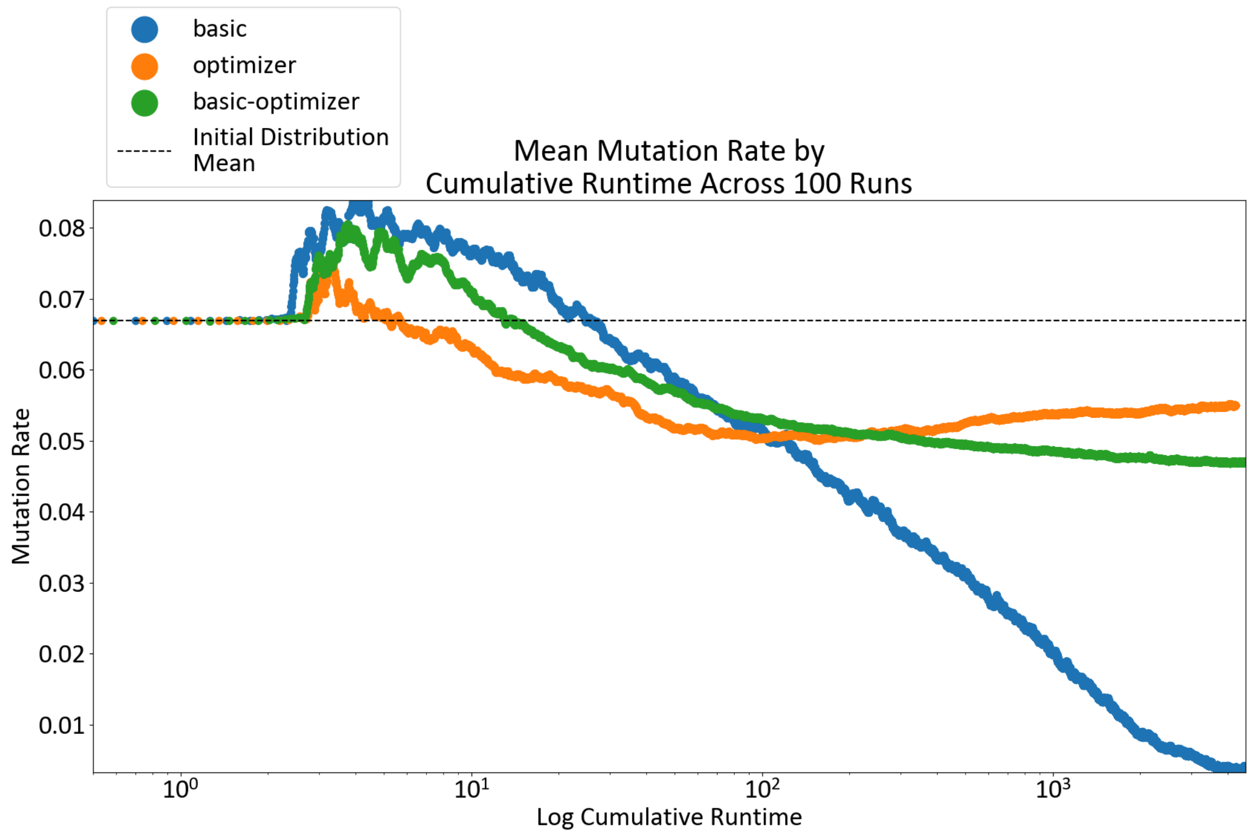

Figure 12: Average Mean of All Mutation Rates, by Breed, by Logarithm Cumulative Runtime

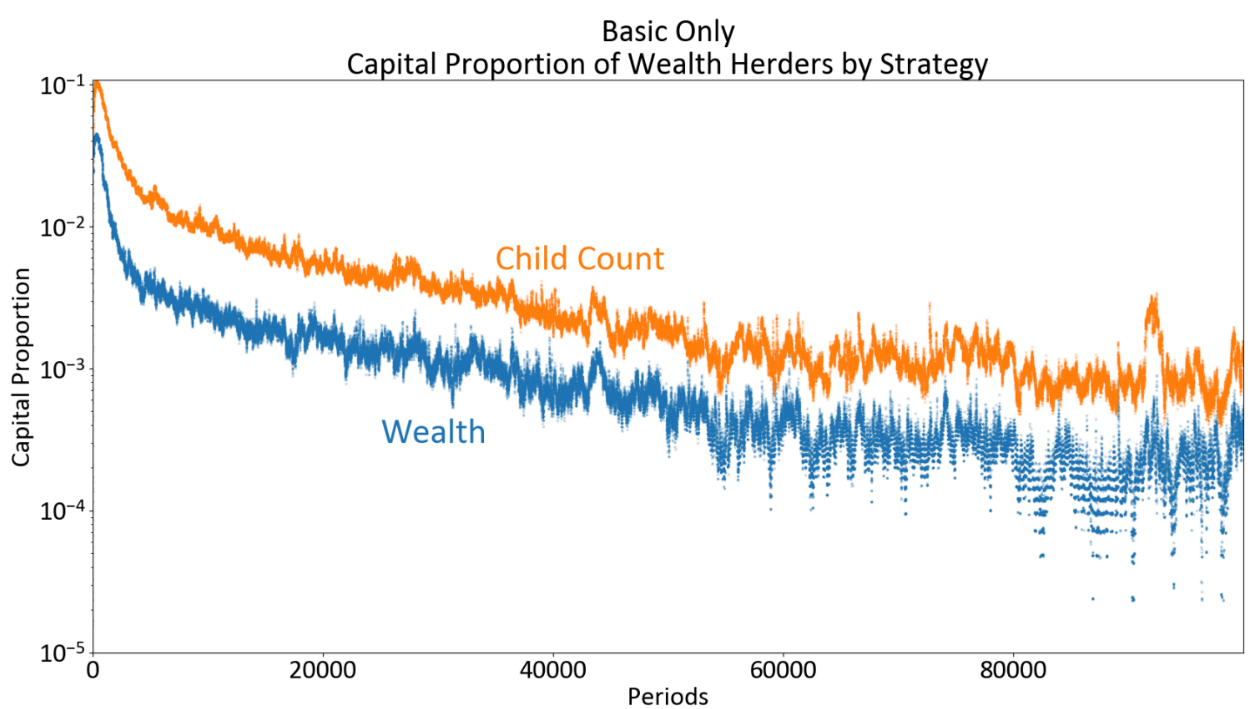

Figure 13: Population Proportion of Herder Strategies by Logarithm of Cumulative Runtime, by Breed Set

As can be seen, the herding strategy of child counting tends to be superior to wealth comparison for the majority of the run, but this advantage diminishes toward the end. A t-test between the averages of all mean observations at each period between the two strategies suggests that the child count strategy tends to be higher throughout most of the run ($\beta = 1.85; t = 3662.76$). A t-test between the final values of each of the two strategies gives a significant estimate ($\beta = 4.69; t = 19.13$). The child counting strategy appears to be relatively more stable – a t-test comparing the absolute values of the logged differences from period to period between the two series shows that the child counting strategy tended to have lower variation ($\beta = 0.91; t = 2558.55$).

Mutation

The average of all agents' mutation rates were tracked each period as the model ran, and each period the population average was taken. One could hypothesize that the mutation rate might, on average, rise in early periods to allow for search of the strategy space, and then fall in later periods once divergences from established norms become counterproductive. However, this model displays mixed results.

***Divergences of** r **from** $E(r)$*

Optimizing agents, only having a single parameter susceptible to mutation, reproduction ratio θ_i, have a mutation rate that drops immediately in runs as they find efficient values for that parameter. The Basic mutation rate acts roughly as hypothesized – it rises early on (though not much) and falls later and only modestly. The combined model acts roughly the same as the Optimizer model, with both higher values early and lower values later.

Figure 12 shows the average mean of all agent's mutation rates, by period, compared to the expected value of the initial distribution of mutation rates. One sample t-tests comparing the average of the 100 final means of the mutation rates to the starting average were performed (Figure 14). The average mutation rate of Basics shows the most significant divergence from its initial mean (t=-108.58), suggesting that Basics use it to shift strategies and learn, as was expected, whereas the Optimizers do not utilize it much for learning given their predefined strategies (t=-5.11). The combined run showed a modest rate of learning (t=-12.34), most of which could have been from the population shift from mostly-Basic to mostly Optimizer, which did not seem to continue after that transition.

The role of mutation in improving fitness can also be conveyed by the final distribution of mutation rates by simulation type. The distribution of the means of the mutation rates in the final period of the simulations, and divergences from their initial distribution, can tell us much about the nature of the evolution of the system.

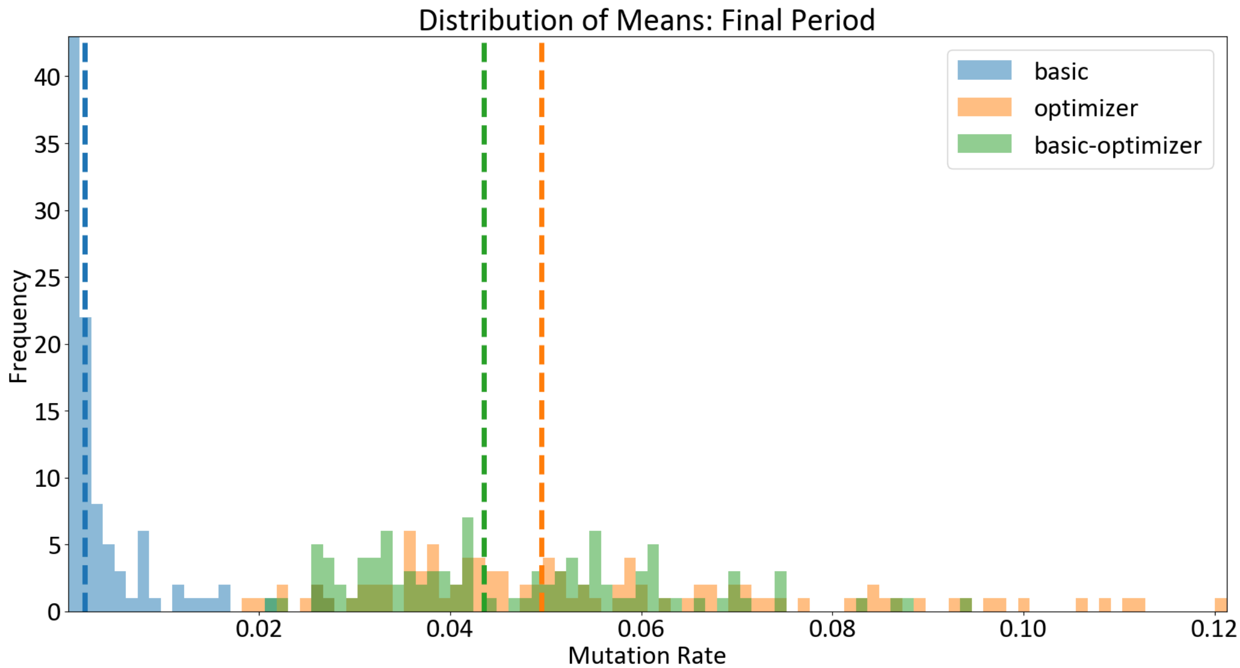

Figure 14: Final Distribution of Means of All Agents Mutation Rates Across Individual Runs

The final distribution of the means of r have resemblance to the reciprocal distribution, with a long right tail. One might expect that values would converge toward a normal distribution since mutation of continuous parameters follows a normal distribution whose mean is the parent attribute value. The skewed distribution from the simulation with only Basic agents is the only distribution with a narrow variance and statistically significant and different mean than the initial distribution. Mutation rates from simulations with only Basic agent mutation rates cluster to values less than 0.02, the minimum value of the initial distribution. Convergence toward low levels of mutation is a means of maintaining the stock of quality knowledge that accumulates across a simulation. This indicates lock-in where Basic agents converge to a limited set of strategies (Arthur 1994). The high level of variance among the means in Optimizer simulations suggests that there is a lack of strong selective pressures with regard to choice of parameter values. Since only one Optimizer attribute (reproduction ratio) is affected by mutation, errors are less costly to the population since selection of an incorrect parameter by a single agent will simply lead to their exit from the simulation. So long as mutation is sufficient to generate the optimal reproduction ratio, the average rate of mutation is not a significant factor affecting survival.

Herding Rates, Mutation Rates and Performance

The mutation rate is a proxy for creativity. High mutation rates across the population should lead to a more diverse set of strategies. However, higher rates of mutation will correlate to higher levels of herders in the population. This creates a buffer on destructive creativity. Higher levels of experimentation require selection of newly created, profitable knowledge. This buffers the costs of experimentation as low quality knowledge will tend to generate lower levels of wealth. A herding agent can shift away from low quality strategies or avoid them altogether. This clear, positive relationship between mutation rates and herding appears across all three sets of simulations (Figure 15).

We reinforce this finding by regressing the percentage of herding agents in a given period on the average rate of mutation in the same period (Figure 16). We fit regressions for a 100 period moving window across all 100 simulations, thus yielding 10,000 observations per regression and a total of 99,900 regressions. Using the same approach, we regress the percentage of herding agents against the level of capital (Figure 17). For both sets of regressions, we visualize the r^2, the beta parameter of the linear term for percent herding agents on the explained variable, and the t-statistic for that estimate. In the plot of the t-statistic, we include a horizontal line at 1.96 to indicate the 5 percent significance level. Overall, we find that correlation between the percentage of herding agents and the average mutation rate is highly and consistently statistically significant. Correlation between the percentage of herding agents and the level of capital is statistically significant for 90,143 of the regressions estimated from the total of 99,900, though estimates of the direction of this relationship are not consistently positive after period 50,000. The r^2 for the same regressions exceeds 0.05 (correlation of 0.224) in 16,335 of these regressions. Of this total, 9269 of those regressions are observed no later than period 40,000.

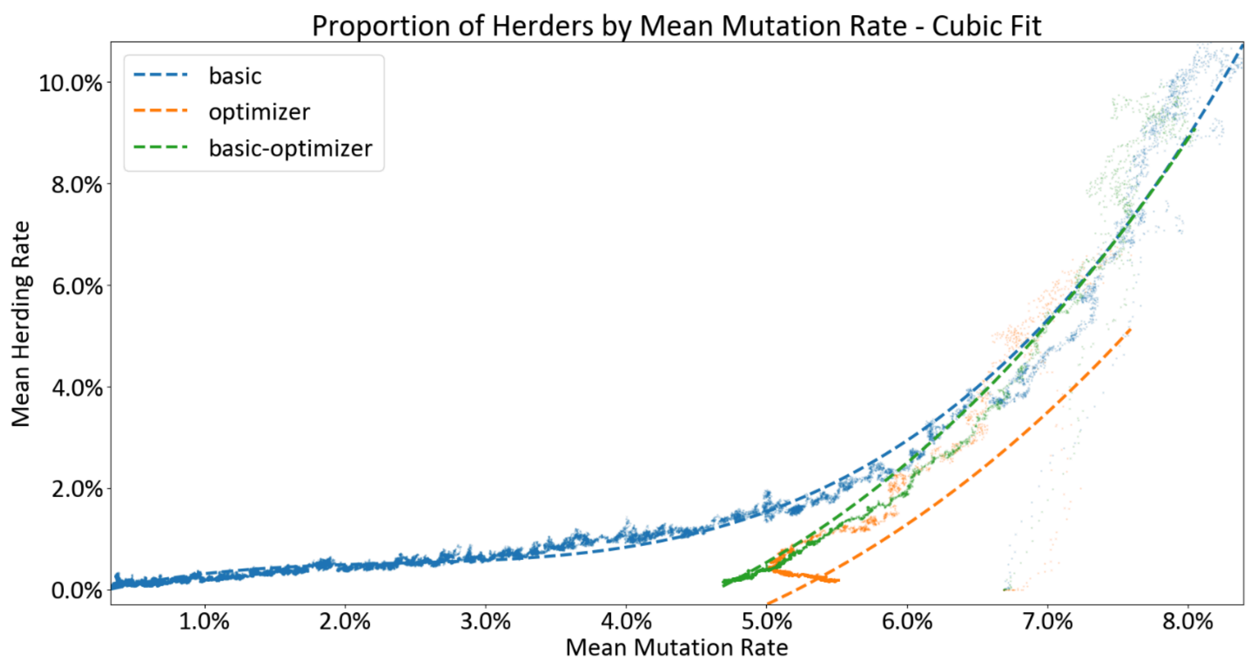

Figure 15: Fitted Cubic Function of for Mean Mutation Rate and Mean Herding Rates (Means taken by period)

Figure 16: Basic Only Simulation. Statistics from Regressions with 100 Period Moving Window, Centered Mean and Median Values over 10000 Regressions.[19]

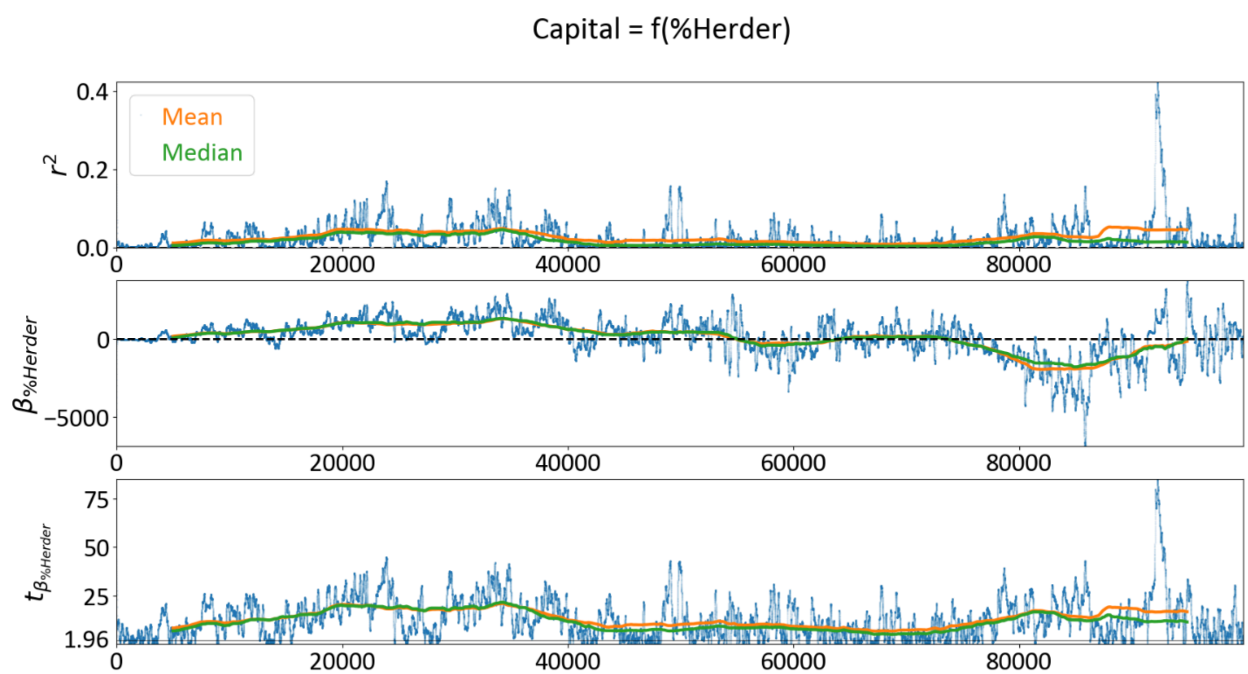

[19] Absolute value of t-values presented for clarity in Figures 16 and 17.

Figure 17: Basic Only Simulation. Statistics from Regressions with 100 Period Moving Window, Centered Mean and Median Values over 10000 Regressions.

Conclusion

We have contributed to literature in computational economics on several margins. First, we have provided a conceptualization of the Market Algorithm. Economic activity closely follows the metaphor of a genetic algorithm whose processes are distributed across economic processes. Of course, real humans are more creative than the robots treated as capital in our model. However, a competitive market of heuristic agents who follow a process of trial and error in developing strategies generates results comparable to a market with agents who maximize CES utility functions. Economists have long-recognized that the perfectly competitive model reasonably describes the outcomes generated by markets, however processes that elaborate competitive outcomes with learning have been more scarce (Axtell and Farmer Forthcoming).[20] The potential diversity of approaches has made general integration of insights from agent-based models less straight-forward compared to models that use the perfectly competitive model as the baseline.

Agents require strategies to constructively determine production and exchange decisions. This includes setting desired quantities and willingness to pay/accept for relevant goods. To interact, agents require a means of identifying relevant trading partners and bargaining with them. Selection mechanisms that remove less efficient strategies as well as those that duplicate more efficient strategies (in part or in whole) allow for dynamic learning that improves productivity much as we observe in modern competitive processes. Agent heterogeneity allows for failure of poorly performing agents and enrichment of high performers with no clear indication ex ante which agents are high performers and which are poor performers.

Second, we have provided a path for macroeconomic analysis of a model with process oriented microfoundations. Macroeconomic variables emerge from these microeconomic processes that are facilitated by meso-level coordination devices (Caton and Wagner 2015; Wagner 2020). Our analysis of macroeconomic variables is informed by new classical analysis, however, the details of agent decision-making processes greatly aid our understanding of the emergent macroeconomic values. By treating the Optimizing agents as our baseline, we expect that the non-Optimizing agents could develop more efficient strategies. Perhaps further sophistication in mutation could aid the search for strategies that move the economy to a higher local maximum or even to a global maximum. It is even possible that these agents could learn to outperform Optimizing agents, though we must leave such a question open for further analysis.

Alchian, A. (1950). "Uncertainty, Evolution, and Economic Theory." *Journal of Political Economy* 58: 211-21.

Alkemade, F., Frenken, K., Hekkert, M. P., and Schwoon, M. (2009). "A Complex Systems Methodology to

[20] This insight traces back at least as far as Frank Knight's presentation of the perfectly competitive model as a juxtaposition to entrepreneurship.

Transition Management." *Journal of Evolutionary Economics* 19: 527-543.

Alkemade, F., Poutre, H. L., Amman, H. M. (2007). "On Social Learning and Robust Evolutionary Algorithm Design in the Cournot Oligopoly Game." *Computational Intelligence* 23 (2): 162-174.

Alkemade, F., Poutre, H. L., Amman, H. M. (2006). "Robust Evolutionary Algorithm Design for Socio-economic Simulation." *Computational Economics* 28: 355-370.

Arrow, K. and Hahn, F. (1971). *General Competitive Analysis*. San Francisco, CA: Holden-Day.

Arthur, W. B. (1994). *Increasing returns and path dependence in the economy*. University of Michigan Press.

Axtell, R. (2016). "Hayek Enriched by Complexity Enriched by Hayek. *Advances in Austrian Economics* 21: 63-121.

Axtell, R. (2005). "The Complexity of Exchange" *Journal of Economic Literature* 115: F193-F210.

Axtell, R. and Farmer, D. (Forthcoming). "Agent-based Modeling in Economics and Finance: Past, Present, and Future." *Journal of Economic Literature*.

Axtell, R. and Guerrero, O. (2019). *Dynamics of Firms from the Bottom Up*. MIT Press

Banerjee, A. V. (1992). A simple model of herd behavior. *The quarterly journal of economics*, *107*(3), 797-817.

Baumol, W. (1968) "Entrepreneurship in Economic Theory" *American Economic Review* 58: 64-71.

Becker, G. (1962). "Rational Action and Economic Theory: A Reply to I. Kirzner." *Journal of Political Economy* 71 (1): 82-83.

Bowles, S., Kirman, A. and Sethi, R. (2017). "Retrospectives: Friedrich Hayek and the Market Algorithm." *The Journal of Economic Perspectives* 31 (3): 215-230.

Bylund, P. (2015). "Signifying Williamson's Contribution to the Transaction Cost Approach: An Agent-Based Simulation of Coasean Transaction Costs and Specialization." *Journal of Management Studies* 52, no. 1: 148–74

Caldwell, B. (2008). *Hayek's Challenge: An Intellectual Biography of F. A. Hayek*. Chicago: University of Chicago Press.

Caton, J. (2017). Entrepreneurship, search costs, and ecological rationality in an agent-based economy. *The Review of Austrian Economics*, *30*, 107-130.

Caton, J. (2019). Creativity in a theory of entrepreneurship. *Journal of Entrepreneurship and Public Policy*, *8*(4), 442-469.

Caton, J. and Wagner, R. (2015). "Volatility in Catallactical Systems: Austrian Cycle Theory Revisited." *Advances in Austrian Economics* 19: 95-117.

Clower, R. and Leijonhufvud, A. (1973 [1981]) “ Say's principle: what it means and doesn't mean” in (Ed.) A. Leijonhfvud. *Information and Coordination* (pp. 79–101). New York: Oxford University Press.

Coase, R. H. (1937). “The Nature of the Firm.” *Economica* 4 (16): 386-405.

Coase, R. (1960). “The Problem of Social Cost.” *The Journal of Law and Economics* 56 (4): 837-877.

Coase, R. H. (1992). “The Institutional Structure of Production.” *American Economic Review* 82 (4): 713-719.

Cobb, C. W., & Douglas, P. H. (1928). “A theory of production.” *The American Economic Review* 18 (1): 139-165.

Dopfer, K., Foster, J., and Potts, J. (2004). “Micro-meso-macro.” *Journal of Evolutionary Economics* 14: 263-279.

Dosi, G. and Roventini, A. (2019). “More is Different . . . and Complex! The Case for Agent-based Macroeconomics.” *Journal of Evolutionary Economics* 29: 1-37.

Easley, D. and López de Prado, M. (2012). “The Volume Clock: Insights into the High-frequency Paradigm.” *The Journal of Portfolio Management* 39 (1): 19-29.

Epstein, J. M., & Axtell, R. (1996). *Growing artificial societies: social science from the bottom up.* Brookings Institution Press.

Fama, E. (1970). “Efficient Capital Markets: A Review of Theory and Empirical Work.” *Journal of Finance*, 25: 383-417.

Felin, T., Kauffman, S., Koppl, R. and Long, G. (2014). “Economic Opportunity and Evolution: Beyond Landscapes and Bounded Rationality.” *Strategic Entrepreneurship Journal,* 8: 269-282.

Foss, N., and P. Klein. 2012. *Organizing Entrepreneurial Judgment: A New Approach to the Firm*. New York: Cambridge University Press.

Gabardo, A. G., Porcile, G., and Pereima, J. B. (2020). “Sectoral Labour Reallocation: An Agent-based Model of Structural Change and Growth.” *Economia* 21 (2): 209-232.

Gigerenzer, G. 2008. *Rationality for Mortals: How People Cope with Uncertainty*. New York: Oxford University Press.

Glasner, D. (2022). “Hayek, Hicks, Radner and Four Equilibrium Concepts: Perfect Foresight, Sequential, Temporary, and Rational Expectations” *Review of Austrian Economics* 35: 39-61.

Gode, D. K. and S. Sunder (1993). “Allocative Efficiency of Markets with Zero-Intelligence Traders: Market as a Partial Substitute for Individual Rationality.” *The Journal of Political Economy,* 101: 119-137.

Gualdi, S., Tarzia, M.,Zamponi, F., and Bouchaud, J. (2015). “Tipping Points in Macroeconomic

Agent-based Models." *Journal of Economic Dynamics and Control* 50: 29-61.

Hayek, F. A. (1933). *Monetary Theory and the Trade Cycle*. New York City: Sentry Press.

Hayek, F.A. (1935a [1948]). "Socialist Calculation I: The Nature and History of the Problem." In *Individualism and Economic Order*. Chicago: University of Chicago Press.

Hayek, F.A. (1935b [1948]). "Socialist Calculation II: The State of the Debate" In *Individualism and Economic Order*. Chicago: University of Chicago Press.

Hayek, F.A. (1935c). *Prices and Production* 2nd Edition. London: Routledge.

Hayek, F.A. (1937). "Economics and Knowledge." Economica 4 (3): 33-54.

Hayek (1940 [1948]). "Socialist Calculation III: The Competitive Solution." In *Individualism and Economic Order*. Chicago: University of Chicago Press.

Hayek (1945). "The Use of Knowledge in Society." *The American Economic Review* 35 (4): 519-530.

Hayek (1948). "The Meaning of Competition." In *Individualism and Economic Order*. Chicago: University of Chicago Press.

Hayek, F. A. (1962) "Rules, Perception, and Intelligibility." in *The Market and Other Orders*. Chicago: University of Chicago Press.

Hayek (1967 [2014]). "Notes on the Evolution of Systems of Rules of Conduct: The Interplay between Rules of Individual Conduct and the Social Order of Actions." In Ed., Caldwell, B., *The Collected Works of F. A. Hayek, Volume 15: The Market and Other Orders*. Chicago: University of Chicago Press.

Hayek (1968 [2002]). "Competition as a Discovery Procedure." *The Quarterly Journal of Austrian Economics* 5 (3): 9-23.

Holian, M., and G. Newell. 2017. "An Agent-Based Model of Entrepreneurship." *Journal of Entrepreneurship and Public Policy* 6 no. 2: 259–70.

Holland, J. (1992). Genetic Algorithms: Computer Programs that 'Evolve' in Ways that Resemble Natural Selection Can Solve Complex Problems Even Their Creators do not Fully Understand. *The Scientific American* 267: 66-72.

Kahneman, D. 2011. *Thinking, Fast and Slow*. New York: Farrar, Straus and Giroux.

Kauffman, S. and Levin, S. (1987). "Toward a General Theory of Adaptive Walks on Rugged Landscapes." *Journal of Theoretical Biology* 128: 11-45.

Kauffman, S. and Johnsen, S. (1991). "Coevolution to the Edge of Chaos: Coupled Fitness Landscapes, Poised States, and Coevolutionary Avalanches." *Journal of Theoretical Biology* 149: 467-505.

Keynes, J.M. (1936 [1964]). *The General Theory of Employment, Interest, and Money*. New York: Harcourt Inc.

Keyhani, M. (2016). “Computer Simulation Studies of the Entrepreneurial Market Process.” In *Complexity in Entrepreneurship and Technology Research: Applications of Emergent and Neglected Methods*, edited by E. Berger and A. Kuckertz, 117–37. Springer.

Keyhani, M. (2019). “Computational Modeling of Entrepreneurship Grounded in Austrian Economics: Insights for Strategic Entrepreneurship and the Opportunity Debate.” *Strategic Entrepreneurship Journal* 13, no. 2: 221–40.

Kirzner, I. (1962). “Rational Action and Economic Theory.” *Journal of Political Economy* 70 (4): 380-385.

Kirzner, I. 1963. *Market Theory and the Price System*. New York: D. Van Nostrand Company.

Lagos, R., & Wright, R. (2005). A unified framework for monetary theory and policy analysis. *Journal of political Economy*, *113*(3): 463-484.

Legnick, M. (2013). “Agent-based Macroeconomics: A Baseline Model.” *Journal of Economic Behavior and Organization*. 86: 102-120.

Lin, W. Y., Lee, W. Y., & Hong, T. P. (2003). Adapting crossover and mutation rates in genetic algorithms. *J. Inf. Sci. Eng.*, *19*(5): 889-903.

Mandelbrot, B. and Hudson, R. L. (2004). *The Misbehavior of Markets: A Fractal View of Financial Turbulence*. New York: Basic Books.

Menger, C. (1871) 2007. *Principles of Economics*. Auburn, AL: Ludwig Von Mises Institute.

Mokyr, J. (2017). *A Culture of Growth: The Origins of the Modern Economy*. Princeton, NJ: Princeton University Press.

Muth, J. (1961). “Rational Expectations and the Theory of Price Movements.” *Econometrica*, 29: 315-335.

Nelson, R. N. and Winter, S. G. (1982). *An Evolutionary Theory of Economic Change*. Cambridge, MA: Belknap Press of Harvard University Press.

Nelson, R. N. and Winter, S. G. (2002). “Evolutionary Theorizing in Economics.” *Journal of Economic Perspectives* 16 (2): 23-46.

Newhard, J. M. (2014). “The Stock Market Speaks: How Dr. Alchian Learned to Build the Bomb.” *Journal of Corporate Finance* 27: 116-132.

Parkes, D. C. and Wellman, M. P. (2015). “Economic Reasoning and Artificial Intelligence.” *Science* 349 (6245): 267-272.

Polanyi, M. (1951). *The Logic of Liberty*. Abingdon, Oxon, UK: Routledge

Schelling, T. (1969). “Models of Segregation.” *The American Economic Review*. 59 (2): 488-493.

Sims, K. (1994). Evolving 3D morphology and behavior by competition. *Artificial life*, 1 (4): 353-372.

Solow, R. M. (1956). “A Contribution to the Theory of Economic Growth.” *Quarterly Journal of Economics* 70: 65-94.

Solow, R. (1957). “Technical Change and the Aggregate Production Function.” *The Review of Economics and Statistics* 39 (3): 312-320.

Stigler, G. and Becker, G. (1977). “De Gustibus non Est Disputandum.” *The American Economic Review* 67: 76-90.

Swan, T. (1956). “Economic Growth and Capital Accumulation.“ *Economic Record* 23 (2): 334-361.

Valente, M. (2014). “An NK-like model for complexity.” *Journal of Evolutionary Economics*. 24: 107-134.

Vriend, N. J. (2001). “A Model of Market-Making.” *European Journal of Economic and Social Systems* 15 (3): 185-202.

Vriend, N. J. (1995). “Self-organization of Markets: An Example of a Computational Approach.” *Computational Economics* 8: 205-231.

Vriend, N. J. (2002). “Was Hayek An ACE?” *Southern Journal of Economics* 68: 811-840.

Wagner, R. E. (2020). *Macroeconomics as Systems Theory: Transcending the Micro-Macro Dichotomy*. Cham, Switzerland: Palgrave Macmillan.

Williamson (1985). *The Economic Institutions of Capitalism*. New York: The Free Press.

Wolff, S. A. and Resnick, R. D. (1987). *Knowledge and Class: A Marxian Critique of Political Economy*. The University of Chicago Press.

Appendix

Total Exchanges

Optimizing agents transact significantly more per period. In the Optimizing agent’s final period (60,000), they had transacted, on average, 160,032,380 times (sum of all exchanges by all agents). This is compared to basics, who had transacted an average of 85,532,300 times by the 60,000th period, and 144,072,800 by their final (100,000th) period. The combined model, in the 60,000th period, had transacted an average of 148,362,020 times (Figure 1A). It will be important to keep this in mind as any measure of the “efficiency” of the market must include time as a cost. The model has no means of penalizing agents for strategies that are computationally costly, but such cost can be accounted for in our analysis using cumulative runtime as a yardstick of development, a variable which mainly diverges across breed types due to differences in total exchanges (Figure 2A).

Although time taken to exchange is comparable, there is a small discrepancy in the slope of runtime against total exchanges. The difference between each slope and level must be explained by other processes occurring in the model. Optimizing agents trade more times per period. Despite the runtime per exchange is greater for Optimizer runs than for basic runs. However, if we compare by level of capital instead of by time period, more trading per unit of capital occurs in simulations with Basic agents (Figure 3A). This may indicate that basic agents require significantly greater information sharing since exchange performed by these less sophisticated agents provide relatively less information than exchanges by Optimizing agents..

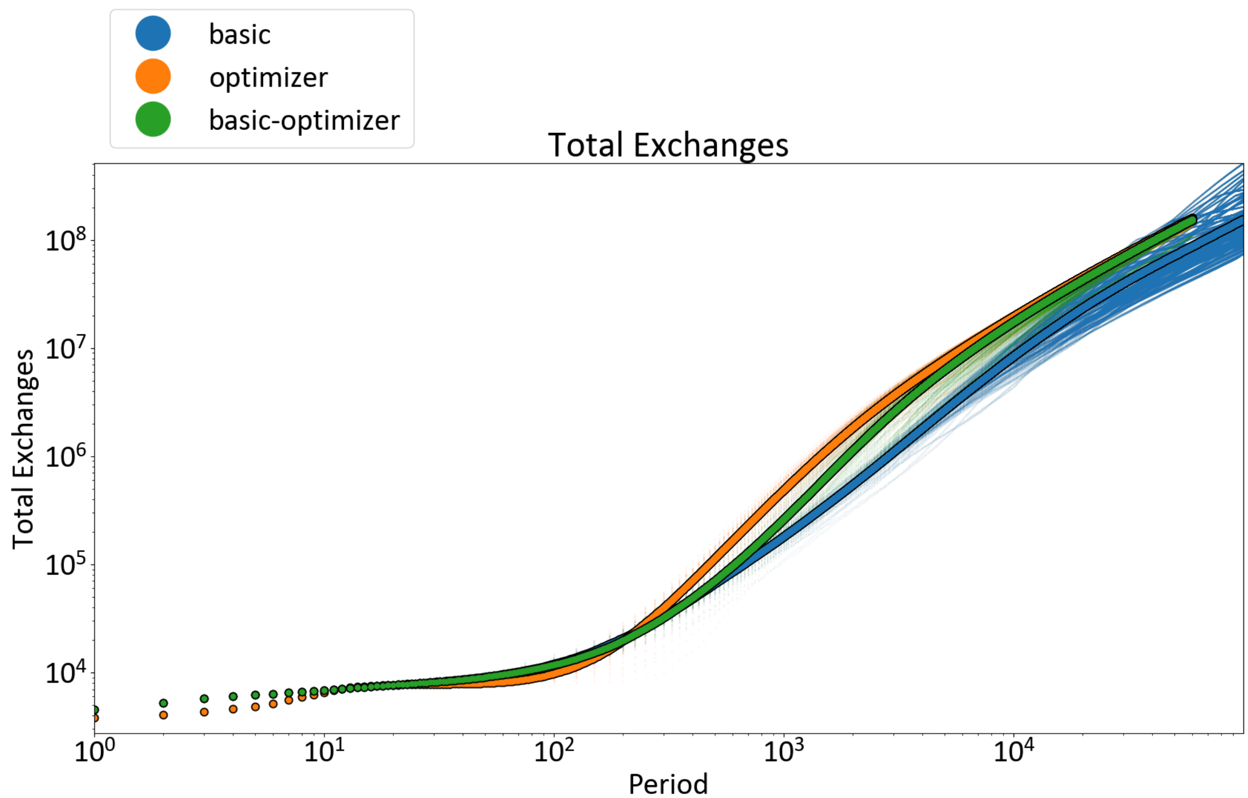

Figure 1A: Total Exchanges by Periods, by Breed Set

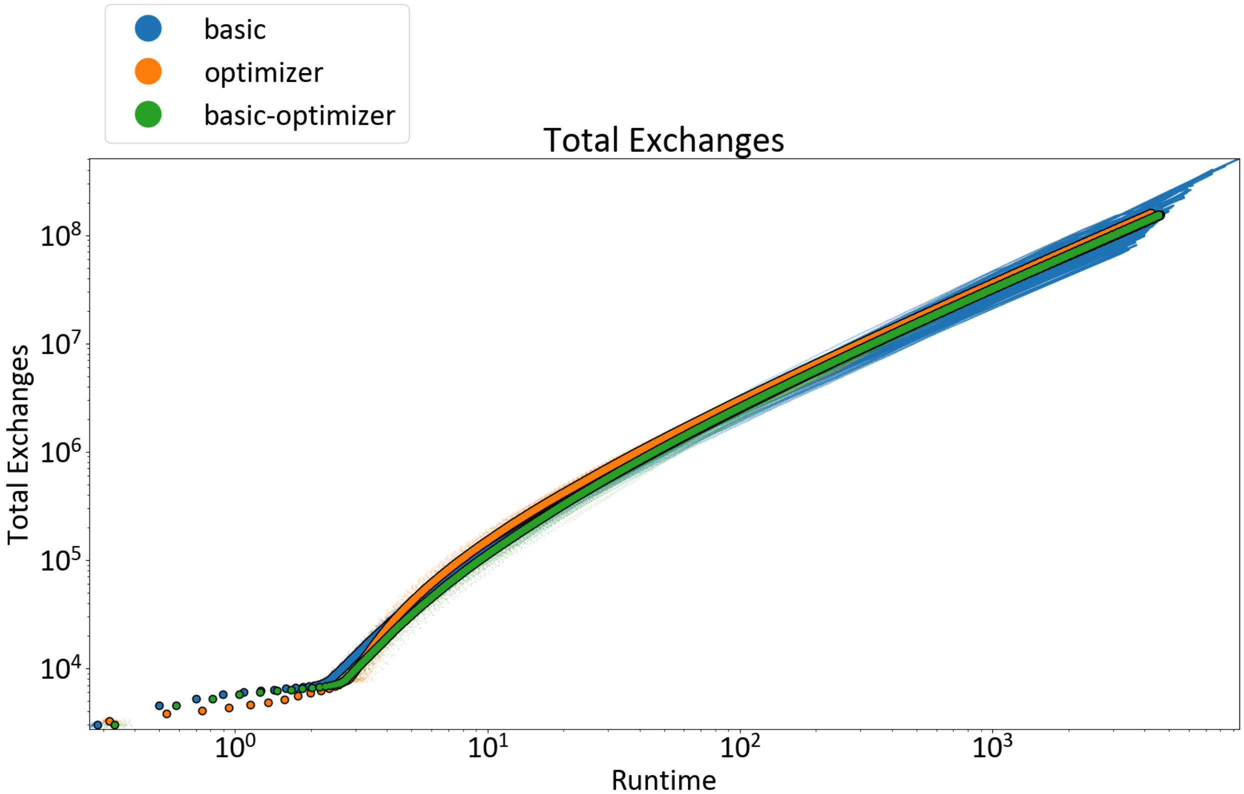

Figure 2A: Log Total Exchanges by Log Cumulative Runtime, by Breed Set

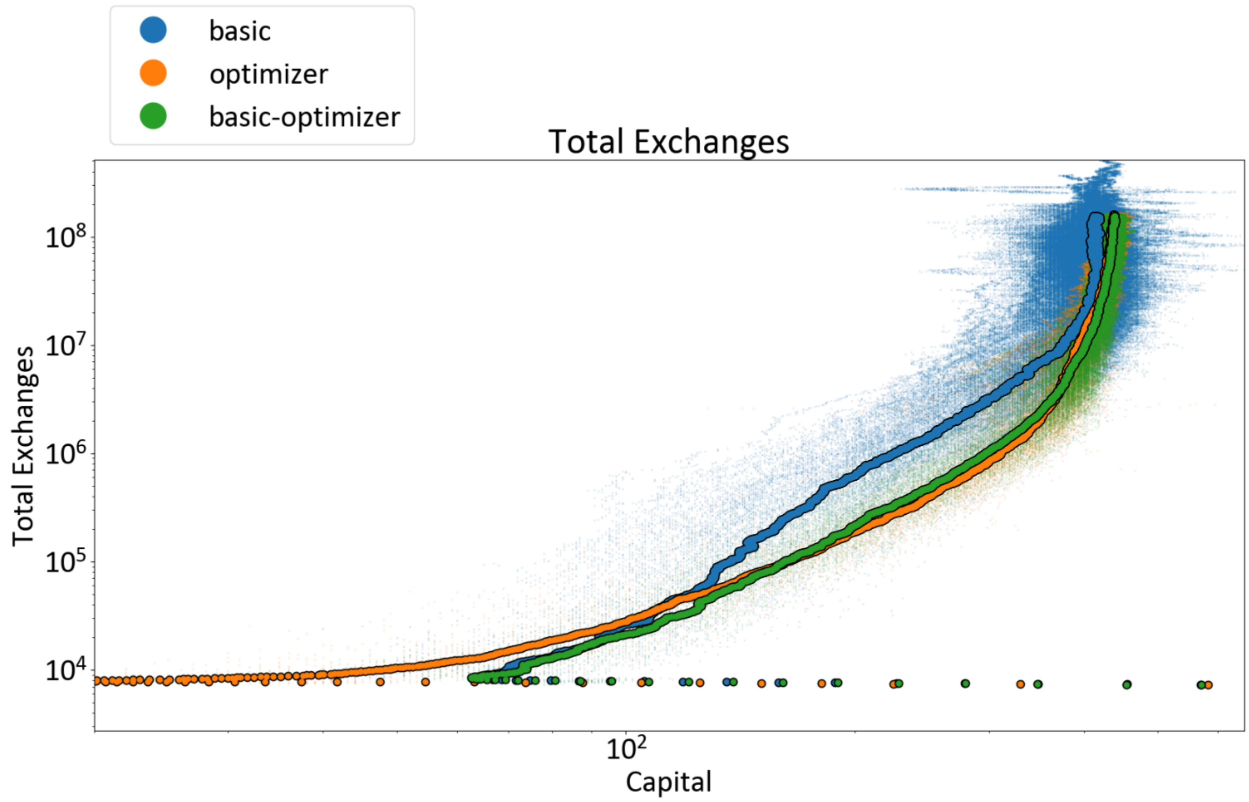

Figure 3A: Log Total Exchanges by Log Capital, by Breed Set

The main advantage of the combined agent runs is the ability to observe how the two breeds interact and compete. These runs allow for mutation which allows an agent to switch from Basic to Optimizer or vice versa if they mutate from their parent's breed or are a herder and discover a wealthier agent whose breed differs from theirs.

Runtime

Optimizers have been shown to transact much more than the basics, presumably due to their more complicated valuation scheme for deciding on worthwhile trades. They also have a more computationally complex valuation scheme for deciding which patch to move to. These increased complexities seem to result in longer computational runtime. For each of the runs for each of the three models, the runtime needed for the period to pass was recorded (Figure 3A).

Figure 6 shows that on a *per period* basis (i.e., time taken for all agents to complete their living tasks), Optimizers take longer to run than the basics. All of the models have increasing runtimes resulting from increasing agent populations, which levels out as they reach their carrying capacities. The combined run converges toward the runtime of the Optimizer-only run as the majority of the agents in these runs become Optimizers.

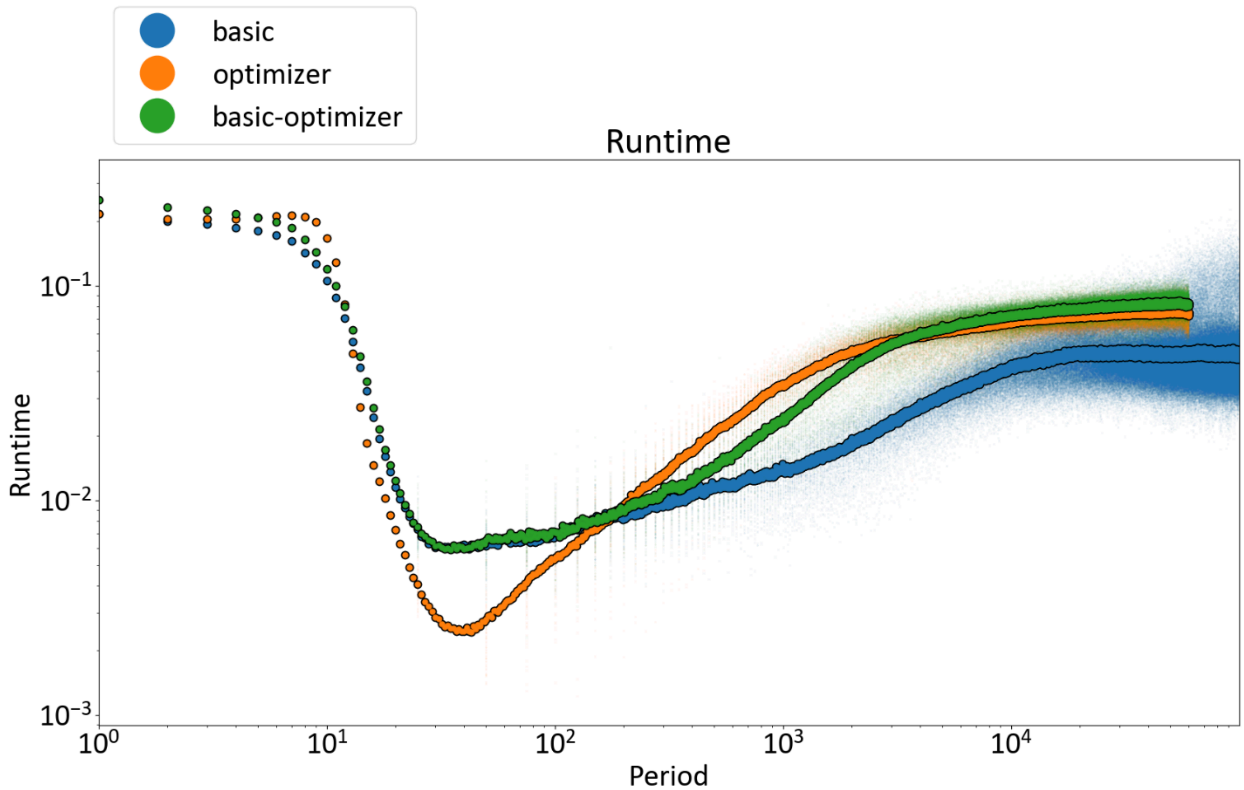

Figure 4A:Log Runtime per Period by Log Period, by Breed Set

We also consider computational runtime on a per period, per agent basis (Figure 4A). For this measure, runtime for the Optimizer and basic-Optimizer (the majority of which are Optimizers) are still significantly higher than that of the basics and increases throughout the run. These results lend credence to the idea that basic agents make simpler decisions than Optimizers. They do not execute the same number of transactions as the Optimizers. The model containing both Optimizers and basics grows larger on a runtime per period per agent basis because Optimizers end up being the predominant breed type in that model. This means that each agent will have an increased number of potential trading partners later in these simulations and will take longer to execute their trading. The Basics, on the other hand, see their wealth per capital stagnate later in runs, as will be shown in a later section. This seems to cause runtime per agent to decrease, as they are no longer increasing their average number of exchanges per trade. The average final total runtime per period for Basics was, on average, 55.8% as high as the same for Optimizers. The average final per-agent runtime for basics was, on average, 60.2% as high as that of Optimizers.