

Implicit Fitness and Heterogeneous Preferences in the Genetic Algorithm

Justin T. H. Smith
University of New Mexico
Department of Economics
New Mexico, USA
1-505-277-1955
jthsmith@unm.edu

ABSTRACT

This paper takes an economic approach to derive an evolutionary learning model based entirely on the endogenous employment of genetic operators in the service of self-interested autonomous agents. Reproductive decisions depend on subjective tradeoffs between the quality and quantity of offspring, avoiding the imposition of an external fitness function as typically used in genetic algorithms in favor of evolving, heterogeneous preferences over reproductive outcomes, expressed via reaction functions. When combined with a density-dependent economic or ecological problem, the implicit fitness approach draws a very different picture of “fitness” than other evolutionary algorithms. An application to learning in a repeated Cournot oligopoly game is developed analytically, predictions tested against a computational simulation. The result is an evolutionarily stable asymmetric equilibrium of much greater average profitability than concentrated Cournot-Nash collusion, while supporting a much larger population of heterogeneous competitors. Implicit-fitness maximizers learn to collude not with their competitors but with their immediate ancestors and descendants to overcome short-run myopia with precommitments through intergenerational wealth transfers.

Categories and Subject Descriptors

I.2.8 [Problem Solving, Control Methods, and Search]:
Heuristic methods; Plan execution, formation, and generation

General Terms

Algorithms, Economics, Experimentation, Theory.

Keywords

genetic algorithms, implicit fitness, utility, preferences, risk aversion, heterogeneity, oligopoly games

1. INTRODUCTION

The Genetic Algorithm (GA) has been a popular paradigm for

evolutionary optimization since its development by Holland [8], successfully applied across a wide range of fields. The idea is simple—evolve a population of solutions to some problem the same way an agriculturalist might improve a strain of corn or breeder might refine a line of dogs. The basic mechanics are simple as well. The GA usually consists of genetic operators—mutation to introduce new variation and cross-over to recombine existing variation—with selection algorithms that determine how this variation is propagated across successive generations. Specifically, under selection, the heritability of some variation is non-random, such that some variants tend to increase in prevalence in the population at the expense of others. This non-random inheritability of continuously-introduced variation is by definition evolution. Over time, the composition of the population changes in the direction favored by selection.

Selection comprises two separate processes. Some members of the population are chosen for reproduction (the remainders are thus selected out), and these breeders are matched in pairs for sexual recombination. The GA is analogous to the artificial selection in that the decisions regarding with whom to mate and how many offspring to produce are made exogenously to meet external criteria. In contrast, evolution via natural selection proceeds from endogenous mate choice and breeding decisions made by heterogeneous individuals in a population, with reproductive behavior itself subject to evolutionary optimization.

This paper takes a fundamentally economic approach to the evolutionary problem. Miller [11] first applied the GA to economics, where it has been used in numerous models of learning and adaptation. Economics has also provided a platform to study the performance of the GA. But economics can also provide the framework for a different treatment of selection in a new evolutionary learning algorithm based on the interaction of subjective beliefs and preferences under density-dependent competition.

The paper proceeds as follows. First, some of the differences that arise between endogenous and exogenous approaches to fitness and reproduction are outlined, examining some of the problems and limitations of selection and learning under the GA. A more dynamic, flexible picture of evolutionary optimization is suggested. Then, an analytical model is developed from which fitness emerges as the implicit result of optimization of individual reproductive decisions. An application to density-dependent competition in a Cournot oligopoly game is developed analytically and implemented in a multi-agent social simulation. An asymmetrical profit-earning steady-state equilibrium is derived

from density dependence, separate from Cournot-Nash and Walrasian outcomes. This and other theoretical predictions are compared to the simulation results. Finally, the paper concludes with a brief discussion of ongoing and future work with this promising new approach to evolutionary optimization through implicit fitness over economic tradeoffs.

2. FITNESS AND UTILITY

2.1 Ex Ante or Ex Post?

The artificial selection approach taken by the GA requires a speculative ranking of all members of the population, such that relative representation in the next generation can be parceled out proportionally. This exogenously-imposed ranking is determined by a so-called “fitness function,” the definition and evaluation of which is often problematic. In this context, “fitness” *determines* reproduction and is measured ex ante.

By contrast, in evolution via natural and sexual selection, endogenous mate choice and offspring investment decisions mean fitness results from the varying ability of agents to produce or capture resources in economic or ecological competition and then use these resources to produce competitive offspring in the face of tradeoffs. *Successful* reproduction determines fitness, properly measured as the geometric mean growth rate (r) [16], only observable ex post.

Essentially, utility is the ex ante estimation of the subjective value of different outcomes. This is precisely what the GA employs a ranking function for, so this will henceforth be referred to as an external or exogenous utility function.

There is no fitness function. The geometric mean growth rate of measured in multi-agent simulations can be compared to theoretical predictions, but by itself r is merely a statistic collected after the fact.

2.2 Heterogeneous Preferences

The GA’s external utility function imposes the same preferences upon every individual in the population. Homogeneity is required to generate a consistent ordered ranking of all individuals, a prerequisite for the application of common selection mechanisms. Roulette-wheel selection takes utility cardinally, assigning a probability for reproduction proportionally over the relative strength of the preference rankings. Tournament selection, on the other hand, makes use of an ordinal utility function, sampling individuals for pure rank tests under this metric. By exogenously specifying a shared utility function of any form, the GA provides a shared set of values. Without such conformity, no games of “how fitter?” or even “who’s fitter?” can be played. This is the primary advantage of the exogenous utility function, but it comes with a price.

Homogeneous utility functions lead to brittle, homomorphic, premature convergence to local optima. This is the inevitable result of any process of learning by imitation, [17]. What results are some of the same issues that plague selective breeders of plants and animals. Reduced genetic and phenotypic diversity follows naturally from the selection algorithm itself. In some pure optimization applications this may not be much of an issue or be readily mitigated, but in multi-agent systems where the emergence of collective behavior is of interest, diversity is key, particularly in complex, open-ended models of social interaction.

A large number of hacks have been developed over the decades to deal with this problem, which reduces replicability of results and can make tuning parameters more of an art than a science. For example, the deep reliance on stochasticity in the GA’s selection algorithm is a tacit admission that the exogenous ranking of individuals is often just wrong. Consider time-variant utility, where fitness depends on the development of many interacting variables through growth, development, and individual learning.

Freeing the utility function from the role of ranking the population allows it to represent individually subjective preferences over both economic and reproductive outcomes related through tradeoffs.

2.3 Density Dependence

The spatial metaphor commonly depicted for the GA is Sewall Wright’s “fitness landscape” [21]. Imagine rolling hills over a space where distances are measured in some genotypic or phenotypic space. Elevations capture cardinal utility preferences, in the ex ante sense of external “fitness” functions. Thus imitative social learning as a hill-climbing algorithm. For generally fixed exogenous utility functions, the landscape is static, affording hill-climbers permanent vacation at the top of local optima.

A more dynamic understanding treats fitness as earned evolutionary profit. Just as economic profit lures economic competition, a positive net evolutionary growth rate ($r > 0$) breeds a larger population and stiffens competition for scarce resources. Evolutionary equilibrium is defined by a zero-profit rule: no net population growth, $r = 0$ as a long run steady-state about which individuals struggle. Spatially, a hill of positive fitness represents a growing population. But the hill erodes away as competition checks growth rates, and in a population crisis may even turn into a sink of net population loss. Evolution never stands still so long as competition continues.

2.4 Population Size

Because the “system” takes care of reproductive decisions in the traditional GA, the population size is generally fixed exogenously or allowed to vary predictably with changes in model inputs. In many models, results can vary significantly depending on such assumptions [1]. The endogenous approach to reproductive behavior depends on the size of simulated populations emerging through interactions between heterogeneous agents. This has its own risks, as the population size may be subject to large swings and even endogenous extinction.

3. A MODEL OF ENDOGENOUS FITNESS

3.1 Evolutionary Optimization

The economic approach to evolutionary optimization was pioneered by Gary Becker, whose work on modeling the demand for children is the basis of the model developed here. Gandolfi et al. [7] offer an easy-to-follow review of much of Becker’s work, synthesizing an economic vision of evolutionary optimization across many different arenas of reproductive and social behavior.

Optimization is about tradeoffs. While models of the tradeoff between the quality and quantity of offspring are nothing new, in economics [4] or biology [15], to this author’s knowledge this is the first time they’ve served as the foundation for a model of evolutionary learning in a social environment.

This approach is also informed by the application of economics to life history theory by Hillard Kaplan, particularly his model of parental investment in embodied capital [10].

3.2 The Problem

The subjective evolutionary objective can be defined as follows. Instead of being directly concerned with their own wealth or health, assume parents maximize some subjective expected utility (EU) over different *reproductive outcomes*:

$$(1) \quad \max_{n,q} EU(n,q;\Pi) = \sum_{R=1}^n pmf(R,n;z(q)) \cdot R \cdot EU(\pi(q)),$$

The choice variables are the quantity n of offspring raised and the quality q of each offspring. These are chosen to maximize the value of the objective function, which depends on the net reproductive rate $R \in [1,n]$ —how many of the optimal $n^* \geq 0$ offspring survive to reproduce given the distribution of mortality risks. The binomial probability mass function $pmf()$ specifies the likelihood of each survivorship outcome R , dependent on the total number of offspring and an assumed independent mortality risk $z \in [0,1]$ which may in turn depend on parental investment $q \geq 0$. The value of each successful reproduction depends on the subjective expected (risk-averse) utility of offspring given some ecological function that turns parental investment into grandchildren, $\pi(q) \geq 0$ over its own probability distribution.

This effectively maximizes the growth rate of the subjective valuation of the optimizing agent's evolutionary lineage, averaged over two generations. If the behavior generated by the optimization is adaptive, it will tend to correlate with positively with geometric mean rate of increase, r —ex post evolutionary fitness.

The expected utility maximization is constrained by the reproductive budget Π :

$$(2) \quad \Pi = (c+q) \cdot n + C,$$

where Π comes from monetary or caloric profits from economic or ecological activity initially funded by the reproducing agent's own endowment of embodied capital. Some fixed cost of building the reproductive system C and marginal cost per offspring c determine upon how many risky offspring n the parent can afford. Assume semelparous reproduction—one litter of size n and then the parent is finished, with every incentive to get the most value from its budget before its generation turns over.

Before attempting to solve the maximization problem, some simplifying assumptions will make for analytical tractability, with an eye toward a simple computational implementation. First, assume subjective certainty about the payoff to parental investment, such that $U(\pi(q)) = EU(\pi(q))$, avoiding the need to contend with nonlinearities over additively separable outcomes. Next, since all offspring are born and raised together and with no means to judge quality before committing to parental investment, assume that it's optimal to invest the same q in each offspring. Thus the expected utility of each offspring is equal, allowing us to pull this scale factor from the cumulative density function. Finally, assume that mortality risks are independent of parental investment and across offspring. Tradeoffs between reproduction and mortality reduction can be added to the basic model, but for

now this yields an extremely useful simplification—only the expected value of the offspring outcome distribution matters, which is well known for the binomial survivorship odds that come from n independent trials at probability of success $1-z$:

$$(3) \quad E(n) = \sum_{R=1}^n pmf(R,n;z) \cdot R = \frac{n}{1-z}.$$

The vastly simplified maximization problem can be expressed with the Lagrangian:

$$(4) \quad \mathfrak{L}(n,q,\lambda;\Pi,c,C) = \frac{n}{1-z} \cdot U(\pi(q)) + \lambda(\Pi - (c+q) \cdot n - C).$$

Inequality constraints on n and q can be considered step-by-step with the Kuhn-Tucker theorem., but note that the $n^* = 0$ corner solution occurs when the parent cannot afford one offspring of even minimal $q = 0$ quality given the fixed and marginal costs of reproduction. Otherwise, the budget constraint binds to a single offspring at suboptimal quality.

Taking the first-order necessary conditions with respect to the choice variables and shadow price λ :

$$(5) \quad \frac{\partial \mathfrak{L}}{\partial n} = 0 = \frac{U(\pi(q))}{1-z} - \lambda^*(c+q^*) \cdot n^*,$$

$$(6) \quad \frac{\partial \mathfrak{L}}{\partial q} = 0 = \frac{n^*}{1-z} \cdot \frac{\partial U(\pi(q^*))}{\partial q} \cdot \frac{\partial \pi(q^*)}{\partial q} - \lambda^* \cdot n^*,$$

$$(7) \quad \frac{\partial \mathfrak{L}}{\partial \lambda} = 0 = \Pi - (c+q^*) \cdot n^* - C$$

Simultaneously solving equations (5) and (6) allows derivation of the optimal parental endowment of capital per offspring:

$$(8) \quad q^* = U(\pi(q^*)) \cdot \left(\frac{\partial U(\pi(q^*))}{\partial q} \cdot \frac{\partial \pi(q^*)}{\partial q} \right)^{-1} - c.$$

Notice that optimal parental investment depends only on the subjective return to capital and the marginal cost of each offspring. The form of the utility function and expected returns function must be known or assumed to analytically solve for the optimal quality.

Critically, since mortality is assumed independent of quality, optimal reproductive behavior does not depend directly on the rate of survivorship. Random death's affect on population dynamics is instead felt indirectly through economic profit Π via the corresponding reduction in competition.

Equation (7) can be rearranged to show that the optimal number of offspring depends on how many of the optimal quality can be afforded within the budget constraint:

$$(9) \quad n^* = \frac{\Pi - C}{(c+q^*)}.$$

Maximizing the sheer number of offspring produced is not an optimal strategy. The tradeoff with quality leads to a reduction in the number of offspring produced to the extent that parents

subjectively value the estimated return to offspring capital investments. The optimum occurs when equal marginal utility is earned from investments in both variables.

3.3 Heterogeneous Preferences

Evolutionary models typically assume risk-neutral maximization of some measure of net reproductive success. Extensive analytical tractability can often be achieved from the use of a linear functional form. And in many cases, the optimal solutions prove invariant to positive monotonic transformation, allowing other conveniences such as taking logs. By assumption, only averages matter. See Charnov [5] for a compact derivation of such a linear evolutionary life-history model inspired by the economic approach. Linear evolutionary models implicitly assume that natural selection will favor expected-value maximizers over risk-averse utility maximizers who subjectively value additional offspring at a diminishing rate. As a result, risk preferences are usually ignored by evolutionary models.

In some problems, including the quality/quantity tradeoff model sketched out above, nonlinear risk preferences do change the subjectively optimal choice variables. The optimal amount of parental investment per offspring in Equation (8) depends almost entirely on risk preferences over the distribution of child economic outcomes. This will be seen more clearly once a specific functional form is assumed.

A central assumption of Becker's approach has been that preferences are stable and similar across populations [17]. In contrast, this model, by letting preferences freely evolve, emphasizes the potential for heterogeneous outcomes.

The dependence of the optimal quality and thus number of offspring on inherited preferences allows the objective function itself to evolve over time. Behavior arises from the interaction of personal preferences and private information in pursuit of self-interested goals. Performance over economic or ecological outcomes is ultimately measured in resultant reproductive outcomes. These linkages are captured in the evolvable reaction functions of Equations (8) and (9).

4. COURNOT OLIGOPOLY GAME

4.1 Background

Consider a simple Cournot competition model, one of the oldest and best-studied problems in game theory. Competing firms produce a homogeneous good for sale in the face of some unknown demand, taking the market price as given. Firms compete on the quantity of output produced. Returns to production are density dependent, as the price is a decreasing function of total production by all firms. A finite population size gives the firms market power, the ability to affect the price via their own output decisions. Firms must simultaneously decide how much to produce, without conspiring with each other. Firms can raise the market price by producing less, but those who don't restrict their output earn more profit at the higher price. At the same time, producing more than other firms lowers the market price, causing negative profits while hurting the over-producer most. In both cases, defecting from the Walrasian zero-profit equilibrium makes a firm relatively worse off than others. On the other hand, there exists a profitable Cournot-Nash equilibrium if all firms learn to coordinate in exploitation of their market power,

colluding to reduce output. The question is, how can firms learn to exploit such profits?

The strategic element of this problem has made it a favorite of economists studying competition and learning for well over a century. With the modern tools of multi-agent simulation, the dynamics of this and other coordination games can be examined in new ways. Such results have shown that a major determinant of convergence to either equilibrium depends on the type of learning algorithm employed. Vriend [20] used the GA in two ways: as it's typically treated as *social learning algorithm* with each firm as a single rule in a population, and second by giving each firm an independent population of rules and its own GA to select among competing beliefs—a hierarchy devised by John Holland called a Classifier System (CS) [9]. The single-rule firms, directly competing with each other in both the economic game and GA, quickly learn to increase output to the Walrasian zero-profit local optimum. In contrast, when the GA operates on rules in the CS, the composition of the multi-ruled firms evolves independently of other firms under an *individual learning algorithm*, and collusion to the profitable Cournot-Nash equilibrium is achieved.

In both cases, the GA was functionally identical, using wealth earned from playing the oligopoly game as fitness with a fixed number of firms. The difference is in how feedback from the economic competition interacts with the evolutionary learning algorithm. The outcome of the Cournot coordination game financially rewards those who defect less from the Walrasian outcome, improving their relative fitness ranking, the so-called "spite effect" [20]. Untethering the reproduction of competing individual rules from economic competition provides a shield from social punishment in a ranked fitness system.

4.2 Investment and Quantity Competition

Firms in the single-stage Cournot model compete on quantity. The embodied capital provided via parental investment comes in the form of a capacity constraint on a firm's output. Thus production is free up to this point, paid for by parental investment. The ultimate tradeoff is between the quantity of offspring produced and the quantity of output each offspring can produce.

Parental investment is a sunk cost, irrelevant to the production decision of each offspring other than providing an upper bound for production with no marginal costs of their own. Because profits are proportional to individual output, incentives ensure production satiates the endowed capacity constraint. These short-run incentives are what drive competition to the Walrasian zero-profit outcome.

Firms can learn to exploit their market power by utilizing their investment in offspring capacity constraints to restrict competition in the next generation. Equilibrium with a positive price is possible because firms can constrain the behavior of future generations for their own long-term genetic benefit.

4.3 Density-Dependent Steady State

The long-run positive-economic-profit equilibrium in the evolutionary game is different from the Cournot-Nash outcome. Consider the density-dependent steady state where the total revenue from the Cournot game equals the amount of capacity supplied onto the market by parents:

$$(10) \quad Q = P(Q) \cdot Q.$$

This reduces to:

$$(11) \quad P(Q) = 1,$$

Which is simply the marginal cost of capital sunk by parents to claim space for their replacement in the next generation. At the density-dependent steady state firms earn one unit of profit for each unit invested by their parent into their endowment. In turn, the next generation's positive Cournot profits are invested into the production capacity of the next generation, so that net evolutionary profits are zero, with no geometrically average population growth, i.e. $r = 0$ on average.

5. COMPUTATIONAL EXAMPLE

5.1 Method

A simple computational implementation of the above economic and evolutionary models was developed to test the feasibility and stability of the evolutionary steady-state. The operation of the simulation is simple. Each generation, each firm produces a quantity of output for sell in the competitive Cournot market. Production is subject to capacity constraints increased by parental investment, with zero marginal cost to the producer as this production has already been paid for. The total market supply is added up and equated with the static demand schedule. The nonnegative price which clears the market is quoted and trades take place, any profit parceled out proportionally to the level of output of each firm.

Then, firms take stock of their profits and make reproductive decisions using their evolved reaction functions. Essentially, reproducing firms endogenously hire genetic operators to create offspring to their own quality specifications, investments in quality taking the form of precommitments to levels of output in the net generation by way of embodied capital reflected in capacity constraints. New offspring make their own decisions as the cycle is repeated indefinitely.

The evolution of the system is determined by the actual mechanisms for forming beliefs and revealing heritable preferences. Details are fleshed out below.

5.2 Beliefs

Preferences serve to evaluate private beliefs about the state of the world. The structure and complexity of these beliefs plays a critical role in the learning method employed. To determine the optimal investment in offspring, an agent needs some way to evaluate the productivity of such investments, $\pi(q^*)$. In the Cournot market, with a common market-clearing price P dependent on total production Q , firm profits are paid as:

$$(12) \quad \Pi = \pi(q_0) = P(Q) \cdot q_0,$$

where q_0 is the capacity constraint chosen by the active firm's parent. Suppose that firms just take a linear extrapolation from the markup they received when producing up to capacity in order to estimate the productivity of parental investment across all q , naively believing their quality choice is unable to move the market price.

$$(13) \quad E(\pi(q)) = \frac{\Pi}{q_0} \cdot q.$$

The expected marginal profit of investment in children's capacity constraints is just the average return the firm's parent received on its own investment. This simple model of adaptive expectations provides a very myopic short-run subjective view of the Cournot market, but it will prove sufficient in the endogenous fitness model. All it requires of firms is the ability to remember the size of their own capacity constraints relative to their profit earned.

5.3 Heterogeneous Preferences

As first proved by von Neumann and Morgenstern [19], a nonlinear utility function can be used cardinally or ordinally to represent preference rankings given certain assumptions about the rationality of decision-making behavior. More complex methods of capturing preference relations can reproduce some of the regular irrationality revealed by behavioral experiments, for example Pin's [13] application of the GA to evolve rankings over simple three-state lotteries in a Machina triangle, but the simplicity and analytical possibilities offered by utility functions make them an obvious first choice. The contrast with the traditional GA is in how the utility functions are used.

Given generally increasing returns to parental investment in evolutionary Cournot competition, preferences must be sufficiently risk-averse to generate the subjectively diminishing returns to investment required for reproduction to be a expected-utility-maximizing strategy. Otherwise the demand for additional quality will never be sated.

It will be necessary to assume a general functional form that can be implemented in a computational model and used to analytically solve the equilibria in question to compare with empirical results. Something that simplifies the math and is easily implemented would be ideal. Suppose utility over outcomes can be specified by a three-parameter power function:

$$(14) \quad U(x) = (k \cdot x + b)^\alpha \geq 0 \\ \forall k \in (0,1], b \in [-100,100], \alpha \in (0,1]$$

Note that $\alpha = 1$ implies a preference for the risk-neutral mean of any distribution, while the power $\alpha \in (0,1)$ can result in either risk aversion or proclivity depending on the offset parameter b and changes in profits, where as b moves from the positive through zero to negative numbers, risk preferences shift from decreasing to constant to increasing relative risk aversion.

A firm's preferences are represented simply by a set of three real-valued numbers in $[0,1)$ appropriately scaled to the bounds reported in Equation (14). These preferences are inherited and subjected to a small probability of point mutations during reproduction, in which case a new preference parameter is randomly generated. This ensures that the entire parameter space can be easily explored to avoid restriction to purely local optima.

The form in Equation (14) is convenient because of the scaling behavior of successive derivatives. Plugging in the simple adaptive expectations belief formation rule, optimal parental investment simplifies to:

$$(15) \quad q^* = \left(\frac{bq_0}{\alpha k \Pi} - c \right) (1 - \alpha^{-1})^{-1}.$$

The optimal amount of investment per offspring depends purely on the three preference parameters, the adaptive expectation of beliefs, and the marginal cost of offspring production. The optimal number of offspring is a slightly messier reaction function:

$$(16) \quad n^* = \frac{\Pi - C}{\left[c + \left(\frac{bq_0}{\alpha k \Pi} - c \right) (1 - \alpha^{-1})^{-1} \right]}.$$

Note that the reaction function is only defined for $\alpha \neq 0$, thus risk-neutral preferences are inconsistent with reproductive demand. Since value in the model comes from the diversification of risk, this is not surprising, but the endogenous fitness approach solidly rejects traditional expected-value maximization.

5.4 Difficulties and Expected Results

Riechmann's [14] comparative study of learning methods' ability to learn collusion in a Cournot-Nash equilibrium depended on informational requirements, the ability to store memories, and the computational and rationality challenges in processing it all. How sophisticated must be the decision-making apparatus? With implicit fitness via subjective expected utility as a new entrant, how should it be expected to perform? Being an individual as opposed to social learning model is a good start, and the algorithm outlined above does remember its own quantity constraint after producing, a memory storage requirement useful to learning to exploit positive profits.

All models that successfully converge on the Cournot-Nash equilibrium required computational complexity equivalent to maximization. The simple imitative learning of the GA is its biggest weakness. In the endogenous fitness approach, given the optimization model above, maximization is implicit even when only using the derived reaction functions for quality and quantity. One place additional complexity is useful is in the handling of non-integer optimal offspring sizes. Simple hacks such as rounding or truncating decimals applies differential and unpredictable harm to firms. A better approach would be to use the subjective expected utility function itself in a direct comparison of the two nearest integer outcomes. This is the method taken in the numerical simulation reported here.

Riechmann's results suggest that information on aggregate demand is a necessary condition for the Cournot-Nash outcome, information that the implicit fitness maximizers lack. As a result, the simple model sketched above is ill-equipped to learning its effect on demand. The prediction of the endogenous fitness model however, is that coordination not with other competitors in the current population but with one's parents and own offspring will be key to achieving positive long-run profits market without such global knowledge. The computational model will serve to test this hypothesis.

The GA with artificial selection fails on all counts above. How does the endogenous fitness model perform? First it must be further parameterized.

5.5 Setting Parameters

Suppose demand is a power law relationship, with parameters chosen such that price is a decreasing function of total market quantity Q :

$$(17) \quad P(Q) = a + m \cdot Q^\beta \quad \text{where} \quad Q = \sum_{i=1}^N q_i.$$

This facilitates easy calculations of total and marginal revenues. On the cost side, zero marginal reproductive costs are problematic in the endogenous fitness approach. They restrict preferences to increasing relative risk-aversion, $b < 0$. Also, a rare mutation causing extreme risk aversion results in a lineage possessed with the desire to diversify itself into a collection of infinitesimal pieces as rapidly as possible. A positive marginal cost to offspring insures against such extreme reproductive behavior and swings in population size. Most importantly, it's the diversion of output capacity and potential profits into reproduction that produces the interesting tradeoffs in the model. No cost to reproduction, no tradeoff. It's the real cost of reproduction that allows a firm to commit the next generation to a lower level of output and potentially lead to positive-profit outcomes.

This dilemma can be solved with careful thought about the fixed cost of building a reproductive system. Consider how integer constraints generally prevent agents from fully maximizing their quality and quantity decisions. Together with satiation of the reproductive budget constraint, this implies that many agents will be induced into producing a single lower-than-optimal offspring. Due to preferences and beliefs, these agents would prefer not to reproduce but to keep all investments in the Cournot market for another turn, and then see if their profits afford them an optimal-quality offspring. Setting the fixed cost equal to the negative of the marginal cost affords agents their first reproduction for free. This fixed cost does not affect the reaction functions. Thus only reproductive decisions that increase the geometric mean growth rate r incur costs.

Simulation parameters are listed in Table 1. Correcting the marginal cost of replacement with a negative fixed cost as mentioned is critical to the emergence of cooperative outcomes. Demand is a simple fixed power law relationship. The initial number of firms with the same capacity constraint is generated, each with random preferences. A small exogenous mortality risk is induced to jog any long-term coalitions that develop and eventually kill any single-firm monopoly to end the simulation,

Table 1. Simulation Parameters

Parameter	Value	Description
C	-10.0	Fixed cost of reproduction
c	10.0	Marginal reproduction cost
a	-1.0	Demand offset
m	5000.0	Demand scale
β	-0.99	Demand power
N_0	20	Initial number of firms
q_0	10.0	Initial quantity constraint

μ	0.01	Mutation rate per pref
z	0.001	Exogenous mortality risk

6. DISCUSSION

6.1 Analytical Predictions

The density-dependent steady state condition is given by Equation (11). Using the demand parameters in Table 1 and the assumed functional forms, the predicted total output level is $Q^{SS} \approx 2,705.6$, well below the competitive Walrasian outcome of $Q^{PC} = 5,449.2$. The evolutionary model is predicted to capture Q^{SS} in profits from the Cournot market and then spending the same on capital investments in production capacity for the next generation, such that the net reproductive rate is zero. The number of firms at this steady-state equilibrium is indeterminate, and thus dependent on the dynamics of competition between heterogeneous preferences:

$$(18) \quad N^{SS} = \beta \left[\frac{1-a}{mQ^{SS}} - 1 \right]^{-1} \quad \text{where} \quad \frac{1-a}{mQ^{SS}} = 1.$$

As such, there's no single Cournot-Nash equilibrium to use a basis for comparison. But it can be shown that the steady-state level of cooperative output is not a Cournot-Nash outcome, which are given by:

$$(19) \quad Q^{CN} = \left[\frac{1-a}{m} \left(\frac{\beta}{N} + 1 \right)^{-1} \right]^{\frac{1}{\beta}}.$$

Meanwhile, the density-dependent steady-state outcome is given by:

$$(20) \quad Q^{SS} = \left[\frac{1-a}{m} \right]^{\frac{1}{\beta}}.$$

For Equation (19) to include Equation (20), $\frac{\beta}{N} = 0$, a contradiction since $\beta < 0$. Interestingly, the number of Cournot-Nash firms required to produce the steady-state level of output can be found by solving the Cournot-Nash equilibrium condition for this amount, as it depends only on demand parameters:

$$(21) \quad N^{CN}(Q^{SS}) \approx \beta \left(\frac{a}{a-1} - 1 \right)^{-1}.$$

For the numerical example in question, $N^{CN}(Q^{SS}) = 1.98$. A Cournot duopoly would not quite be able to maintain the joint profits achieved in the asymmetrical density-dependent steady-state. Does this mean the evolutionary steady-state only supports a low number of firms? The simulation model can help resolve the issues.

6.2 Computational Results

Computation experiments are in their infancy on this project, but initial results can be briefly described. With the above parameterization, the model rapidly and invariably converges to the density-dependent steady state earning positive economic profit as predicted. This is not surprising given the assumptions made above, as intergeneration resource transfers must pay for themselves in the repeated equilibrium.

What is intriguing in achieving such a flow of profits is that they're not monopolized by a small number of firms taking advantage of increasing returns to scale. Under riskier environments, with high exogenous mortality rates, the noise is more easily exploited by larger competitors and lock-in occurs, but under milder conditions these high-profit steady-states support a large number of competitors.

Even though there's hardly any mortality risk, a tenth of a percent chance per generation, heterogeneous preferences induce firms to cooperatively withhold enough production to raise the market price significantly above each generation's own marginal production costs, extracting more profits from the Cournot market than a cooperating duopoly.

With asexual reproduction and near 100% heritability this effectively creates an overlapping generations model where each "generation" could be considered a time-specific manifestation of a single agent, preferences fixed but for rare mutations while beliefs are updated over multiple periods until accumulating enough profits to occasionally diversify through reproduction, $n \geq 2$. Parental investments in the form of precommitted quantity constraints allow these overlapping selves to overcome the time inconsistency problems of their myopic short-run expectations, restricting output to produce higher per-firm profits than the Cournot-Nash equilibrium. Long-run cross-generational evolutionary profits still converge to zero with zero net population growth.

The relationship between fixed and marginal costs supports the evolution of density-dependent populations in the Cournot market. This is because parental investment is constant returns to scale in the no-growth case when the effective marginal cost is zero. If returns to scale are increasing at the long-run evolutionary equilibrium ($r=0$), the market will be entirely captured by a single monopolist. Using the fixed cost to negate the positive externality caused by a marginal cost for the first reproduction assures that increasing returns only kick in for $n \geq 2$, allowing the emergence of economically profitable multi-firm oligopolies at the density-dependent steady state.

7. CONCLUSION

With endogenous fitness, every agent is effectively equipped with its own set of genetic operators at the disposal of its subjective reproductive preferences, constrained by the agent's ability to acquire profits from economic or ecological competition. The price a firm pays to use these genetic operators is endogenously determined by the firm itself, reflecting the subjective value reproduction yields from decreased exposure to evolutionary risk. Selection is natural in that agents who fail to compete and lineages that fail to spread are removed from the population over time. But the genes of different agents are not in direct competition as in a traditional GA; instead they compete indirectly through their effects on behavior. The result is an *individual evolutionary learning algorithm* that stands apart from existing methods.

As this work is in its preliminary stages, much needs to be done to analytically and experimentally explore the parameter space and flesh out the character of the density-dependent steady state. In addition, many obvious extensions suggest themselves. First it's natural to want to explicitly model probabilistic economic

outcomes to study the link between reproductive and economic risk behaviors. A fuller overlapping generations approach would allow the incorporation of a larger host of economic and biological tradeoffs, including growth versus investments in improved survivorship [6]. Sexual reproduction with endogenous evolving mate choice algorithms can be derived from the work of Becker and others. Discounting by subjective temporal preferences as well as relatedness for inclusive fitness is possible. The options are endless, but with each extension analytical solutions become more difficult. Still, the specification of the reaction functions is simpler than developing good explicit utility functions for complex tasks. And as derivations become intractable the methodology of directly evolving non-parametric reaction functions rather than deriving structural forms from underlying optimization models may prove fruitful.

The intersection of economic and evolutionary approaches to complex dynamic problems has proven very rich over the last several decades, a trend that seems to be growing as multi-agent models with learning and complex social dynamics are applied to more and more problems. The extension of economic preferences to evolutionary outcomes can lead to a large array of implicit fitness models evolving in the face of tradeoffs.

8. REFERENCES

- [1] Alam, Shah Jamal, Edmonds, B., and Meyer, R. Identifying Structural Changes in Networks Generated from Agent-Based Social Simulation Models. *Lecture Notes in Computer Science*, 5044 (2009), 298-307.
- [2] Alkemade, Floortje, La Poutré, H., and Amman, H. M. Robust Evolutionary Algorithm Design for socio-economic simulation. *Computational Economics*, 28 (2006), 355-370.
- [3] Arifovic, Jasmina, and Maschek, Michael K. Revisiting individual evolutionary learning in the cobweb model—an illustration of the virtual spite-effect. *Computational Economics*, 28 (2006), 333-354.
- [4] Becker, Gary S. and Lewis, H. G. On the interaction between the quantity and quality of children. *Journal of Political Economy*, 81:2 II (1973), 279-288.
- [5] Charnov, Eric L. Trade-off-invariant rules for evolutionarily stable life histories. *Nature*, 387:22 (1997), 393-394.
- [6] Charnov, Eric L. The optimal balance between growth rate and survival in mammals. *Evolutionary Ecology Research*, 6 (2004), 307-313.
- [7] Gandolfi, Arthur E., Gandolfi, A. S., and Barash, D.P. *Economics as an Evolutionary Science*. Transaction Publishers, (2002).
- [8] Holland, John H. *Adaptation in Natural and Artificial Systems*. University of Michigan Press, (1977).
- [9] Holland, John H. and Reitman, J.S. Cognitive systems based on adaptive algorithms. In *Pattern Directed Inference Systems*, Academic Press, (1978).
- [10] Kaplan, Hillard S. A theory of fertility and parental investment in traditional and modern societies. *Yearbook of Physical Anthropology*, 39 (1996), 91-135.
- [11] Miller, John H. A genetic model of adaptive economic behavior. University of Michigan working paper, (1986).
- [12] Mueller, L. D, Guo, P. Z, and Ayala, F. J. Density-dependent natural selection and trade-offs in life history traits. *Science*, 253:5018 (1991), 433-435.
- [13] Pin, Paolo. Evolution of Risk Preferences. *Mathematical Methods in Economics and Finance*, 1:1 (2006), 65-76.
- [14] Riechmann, Thomas. Cournot or Walras? Long run results in oligopoly games. *Journal of Institutional and Theoretical Economics*, 162:4, (2006), 702-720.
- [15] Smith, Christopher C. and Fretwell, S.D. The optimal balance between size and number of offspring. *The American Naturalist*, 108: 962 (1974), 499-506.
- [16] Stearns, Stephen C. Daniel Bernoulli (1738): Evolution and economics under risk. *Journal of Biosciences*, 25:3, (2000), 221-228.
- [17] Stigler, George J. and Becker, Gary S. DeGustibus non est disputandum. *American Economic Review*, (March 1977).
- [18] Vega-Redondo, Vega. The evolution of Walrasian behavior. *Econometrica*, 65:2 (1997), 375-384.
- [19] Von Neumann, John and Morgenstern, Oskar. *Theory of Games and Economic Behavior*, Princeton University. (1944).
- [20] Vriend, Nicolaas J. An illustration of the essential difference between individual and social learning, and its consequences for computational analyses. *Journal of Economic Dynamics and Control*, 24 (2000), 1-19.
- [21] Wright, Sewall. The distribution of gene differences in populations. *Proceedings of the National Academy of Sciences*. 23:6 (1937), 307-320.