

On the Evolution of Investment Strategies and the Kelly Rule—A Darwinian Approach*

TERJE LENSBERG¹ and KLAUS REINER SCHENK-HOPPÉ²

¹Norwegian School of Economics and Business Administration; ²University of Leeds

Abstract. This paper complements theoretical studies on the Kelly rule in evolutionary finance by studying a Darwinian model of selection and reproduction in which the diversity of investment strategies is maintained through genetic programming. We find that investment strategies which optimize long-term performance can emerge in markets populated by unsophisticated investors. Regardless whether the market is complete or incomplete and whether states are i.i.d. or Markov, the Kelly rule is obtained as the asymptotic outcome. With price-dependent rather than just state-dependent investment strategies, the market portfolio plays an important role as a protection against severe losses in volatile markets.

JEL Classification: G11, G12, C63

1. Introduction

In this paper, we pursue a Darwinian approach to the study of the evolution of investment strategies in financial markets with short-lived assets. The model comprises the two main processes, selection and reproduction, in a genetic programming framework. According to this approach, center stage is occupied by the population which embodies the investment skills of many individual strategies. Our investors are simple-minded and unsophisticated in the sense that they follow preprogrammed behavior rules which are the result of mutations and crossovers. This simplicity is a key factor, as it opens up the possibility of studying, in quite a realistic context, the validity of equilibrium predictions derived from theoretical models that impose strong assumptions on the market dynamics, as well as on individuals' rationality or learning behavior. Our approach complements these models by replacing their rationality assumptions with a Darwinian selection mechanism in

* We are most grateful for helpful and constructive comments from an anonymous referee and from the editor, Peter Bossaerts. Thanks are also due to Jarle Møen for a helpful discussion. Financial support from Leeds University Business School and the National Centre of Competence in Research "Financial Valuation and Risk Management" is gratefully acknowledged.

which investment strategies emerge with a degree of risk aversion that is appropriate for survival. This approach also provides information on the stability properties of equilibria.

In the absence of a reproductive process that creates diversity, the market selection dynamics for short-lived assets are well-studied from an evolutionary perspective (Amir et al., 2005; Evstigneev et al., 2002 and 2006), as well as from a Bayesian viewpoint (Blume and Easley, 1992). The selection pressure in the model considered here is provided by the wealth dynamics which give investment strategies that accumulate more wealth than others a stronger impact on market prices and allocations. Kelly (1956) proved that if markets are complete and consist only of Arrow securities, the rule of "betting one's beliefs" eventually accumulates all wealth. This rule prescribes dividing wealth across assets according to the probability of their paying off. In incomplete markets, the Kelly rule generalizes to the strategy of setting portfolio weights equal to the assets' expected relative payoffs (Hens and Schenk-Hoppé, 2005). The Kelly rule always has a non-negative growth rate relative to the market, and its growth rate is strictly positive if the other investors do not hold the market portfolio, which yields zero relative growth. Moreover, if relative asset prices are given by the portfolio weights of the Kelly rule, this strategy corresponds to the log-optimal investment (see Algoet and Cover, 1988; Hens and Schenk-Hoppé, 2005). In summary, these papers find that the Kelly rule is selected by the market, in the sense of accumulating total market wealth in the long run. However, these results hinge on restrictive assumptions: either there is an investor who follows the Kelly rule from the beginning, or there is a Bayesian learner with logarithmic preferences whose prior includes the true model of the economy.

Related studies on market selection within general equilibrium models provide less clear-cut results. In dynamically complete markets, the Bayesian learner with the most accurate beliefs prevails in the long run, regardless of risk preferences (Blume and Easley, 2006; Sandroni, 2000). Since only beliefs matter for survival in complete markets, the Kelly rule does not play any prominent role in this model. For incomplete markets, even the link between the accuracy of beliefs and survival does not hold in general. Blume and Easley (2006) provide an example in which an agent with wrong beliefs drives out a trader with correct beliefs, even though the latter maximizes the logarithmic growth rate of wealth. The agent with incorrect beliefs judges returns too optimistically in relation to the true probability measure and "outsaves" the agent with correct beliefs. In complete markets, Pareto optimality of the equilibrium allocation precludes this from happening.

Research on the performance of rational versus irrational traders, as well as the price impact of the latter, is also closely related to this paper because traders are characterized by rules of behavior. De Long et al. (1990 and 1992) show that noise traders can survive in markets and impact the prices in the long run. Survival of a trader, however, is not a necessary condition for price

impact, as clearly demonstrated by Kogan et al. (2006). They find that, even if a trader's wealth tends to zero, he can influence prices in the long run through his impact on the state-price density in states with low payoffs.

None of these results can be regarded as satisfactory from an applied point of view. First, the general equilibrium approach to market selection does not offer robust results on asset prices and their long-run dynamics. Since a trader with an arbitrary (standard) utility function can dominate the market in the long term, prices are not pinned down. Second, neither the general equilibrium approach nor the existing literature on the Kelly rule allows for the entry of new investment strategies or the exit of unsuccessful ones. Third, one cannot construct the Kelly rule without sufficient information about asset payoffs and underlying probabilities. Fourth, without this information it is not possible to construct an appropriate prior and Bayesian learning might fail. Finally, the chances of seeing the Kelly rule emerge in a market in which the traders do not have any knowledge of the theoretical results might be slim.

The approach used here provides an evolutionary finance model in which these shortcomings are overcome. We allow the set of investment strategies to develop over time in a Darwinian fashion, using a genetic programming algorithm (Koza, 1992; Smith, 1980) to model the evolutionary process. Investment strategies are represented as computer programs, and new investment strategies are produced by genetic recombination of strategies that have performed well in the past. This maintains diversity in the pool of investment strategies and occasionally produces new strategies with superior capability to generate wealth for those investors who adopt them.

Genetic programming belongs to a tradition in computer science which employs the principles of Darwinian evolution to breed artificially intelligent agents who can solve complex tasks (Holland, 1975). A pioneering application in finance is Arifovic (1996)'s analysis of exchange rate fluctuations in an overlapping generations economy. Following Neely et al. (1997) and Allen and Karjalainen (1999), there is a growing literature on the usage of genetic algorithms in identifying profitable trading rules in financial markets. In a related paper, Lensberg (1999) analyzes a special case of Blume and Easley (1992)'s investment model in a genetic programming framework and finds that surviving behavior rules indeed act as expected log-utility maximizers with Bayesian updating.

Our model describes a situation where investors progressively improve their skills without Bayesian rationality or other sophisticated learning procedures. This is possible because investor skill is tacit knowledge, acquired through imitation and repeated trial-and-error, as emphasized by Polanyi (1967). Successful investors in our framework are like Friedman (1953)'s billiard players who somehow manage to get it right, but are unable to explain how.

Market clearing is modeled as in the market game of Shapley and Shubik (1977). This allows us to focus on the strategic aspects of portfolio

choice while abstracting from details of implementation. Each agent has an investment strategy which is used to select portfolio weights. These weights may depend on information about the current state of the economy and on historical price information, but the agents must make their decision without definite knowledge about the prices that prevail when their decision is implemented in the market. Short selling is not permitted. This excludes one influence that is capable of correcting prices, although, as shown by De Long et al. (1990 and 1991), this does not necessarily occur. In each new period, each agent carries over a portfolio of asset holdings from the previous period and receives state contingent dividends. The dividends are reinvested in fresh assets, each of which becomes available in a fixed unit supply. The Shapley-Shubik mechanism clears the market by simply equating the market capitalization of each asset (its price) to the total wealth invested in it.

All of this is—by and large—faithful to the original setting proposed by Kelly (1956), and yet it is in contrast to the general equilibrium approach. However, there is a close relation to the latter. Market game equilibria converge to competitive equilibria as the number of traders increases (Shapley and Shubik, 1977), and any equilibrium price sequence can be mimicked within our framework when the portfolio weights are allowed to be time- and history-dependent (Amir et al., 2005).

We perform four sets of experiments in which genetic programs are supplied only with information about the current state of the economy. The experiments differ with respect to the market structure and the fundamental stochastic process. The findings are positive throughout. The Kelly rule emerges from the population of genetic programs in all experiments. Throughout a transitory period, however, the market is closer to the theoretical equilibrium than are the individual investment strategies, an observation that is in line with experimental findings (see Bossaerts et al., 2005). The information spreads efficiently through the population so that, eventually, the majority of the population follows the Kelly rule, i.e., convergence of market and strategies prevails. In this sense the optimal investment strategy is indeed learned. The findings are robust with respect to the level of noise that is generated by mutation, crossovers and the inflow of wealth.

The impact of the availability of price information is studied in detail within one of these experiments. In that setting, the Kelly rule coincides with the growth-optimal investment, which is independent of market prices even though it depends on the state. From a genetic programming perspective, the problem becomes considerably harder to solve because the number of inputs to the decision problem increases from 2 (state and asset) to 5 (state, asset and three prices). Market prices and strategies still converge to those predicted by the Kelly rule, but the transitional dynamics and successful investment styles are qualitatively different. An analysis of wealthy genetic programs shows that the market portfolio plays an important role. Purchasing the market

portfolio ensures that a constant share of the total wealth is maintained and low levels of wealth are avoided. This increases a strategy's chance of survival by reducing the probability of being deleted by the genetic recombination process. If state-contingent market clearing prices are perfectly anticipated, such a strategy neither loses nor gains in the market dynamics. In volatile markets, there is slippage, but even an approximate market portfolio provides considerable downside protection.

This study is numerical, and one might ask whether analytical results exist that would put more solid ground under our findings. From a mathematical perspective, the major challenge—and also the main departure from the model considered in Amir et al. (2005)—arises from endowing newly created strategies with wealth. This process depends on the state of the market rather than being an exogenous and purely stochastic event. Given that the analysis in Amir et al. (2005) is quite sophisticated, analytical results do not appear to be straightforward. Moreover, Amir et al. (2005) and related papers by Blume and Easley (1992) and Evstigneev et al. (2002) only provide results on the asymptotic dynamics of the basic model. They neither study transient behavior nor provide estimates on the speed of convergence for the short- and medium-term. Our paper, however, is particularly concerned with these two issues. Particular emphasis is placed on the properties of the wealthiest investment strategy where little or nothing is known, in general, because of its highly path-dependent nature.

The remainder of the paper is organized as follows. Section 2 presents the Darwinian model of a financial market in three steps: wealth dynamics, 2.1, investment strategy dynamics, 2.2, and implementation of the genetic algorithm, 2.3. The experiments are presented in Section 3: General results are given in 3.1, and a detailed analysis of the complete market with Markov states is provided in 3.2. The case of price-dependent strategies is considered in Section 4. Section 5 concludes.

2. A Darwinian Finance Model

This section introduces an evolutionary model of a financial market with short-lived assets. The model incorporates the two Darwinian processes of selection and reproduction. Selection is given by the wealth dynamics among a fixed set of investment strategies, and reproduction imposes a dynamic structure on the set of strategies itself. Both processes are captured here by implementing the model using genetic programming.

2.1 WEALTH DYNAMICS

We briefly recall the evolutionary finance model studied in Amir et al. (2005), Evstigneev et al. (2002 and 2006) and Hens and Schenk-Hoppé (2005). This

model describes the wealth dynamics among a given pool of I investment strategies that interact in a financial market. Time is discrete, $t = 0, 1, 2, \dots$. There are K assets with random payoffs $D^k(s) \geq 0$, $k = 1, \dots, K$, with $\sum_{k=1}^K D^k(s) > 0$. Here $s = 1, \dots, S$ denotes the state of nature. Each asset is short-lived and in fixed supply of one unit. An investment strategy is a sequence of time- and history-dependent vectors of portfolio weights $\lambda_t^i = (\lambda_{1,t}^i, \dots, \lambda_{K,t}^i)$, $\lambda_{k,t}^i \geq 0$ and $\sum_{k=1}^K \lambda_{k,t}^i = 1$.

The price of asset k at time t is given by $q_{k,t} := \lambda_{k,t} w_t = \sum_{i=1}^I \lambda_{k,t}^i w_t^i$, i.e., $q_{k,t}$ is equal to the total amount of wealth invested in asset k . Investor i 's portfolio holdings in asset k is determined as $\lambda_{k,t}^i w_t^i / \lambda_{k,t} w_t = \lambda_{k,t}^i w_t^i / q_{k,t}$ which is a fraction of the one unit supplied. Each strategy's wealth in the next period is determined by the total receipts of random asset payoffs which are distributed according to the portfolio holdings.

The evolution of the distribution of wealth $w_t = (w_t^1, \dots, w_t^I)$ across investment strategies is governed by

$$w_{t+1}^i = \sum_{k=1}^K D^k(s_{t+1}) \frac{\lambda_{k,t}^i w_t^i}{\lambda_{k,t} w_t} \quad (1)$$

for $i = 1, \dots, I$. The state s_{t+1} is randomly drawn according to a given probability distribution.

The pricing equation $q_{k,t} = \lambda_{k,t} w_t$ merits a more detailed discussion. It is the market clearing condition of the Shapley and Shubik (1977) market game, which equates the market capitalization of an asset to the total amount of wealth invested in it. Since each asset is in one-unit fixed supply, this is simply the price of asset k . Thus, the Shapley-Shubik market game simultaneously clears, with each time step, K markets and yields a unique short-term equilibrium price vector. That is quite different from agent-based models, where usually only one market-clearing price is needed (Hommes, 2001). From the definition, it is clear that each strategy's impact on the price is proportional to its wealth. Short selling is excluded to avoid bankruptcy, which would be prevalent in the presence of a short-run equilibrium. For Arrow securities, which have a positive payoff in one state of the world and pay zero otherwise, the price determines the odds of the corresponding bet; they are given by $D^k(s_{t+1} = k) / q_{k,t}$. In this respect, the market-clearing mechanism corresponds to the one used in parimutuel betting markets.

Equation (1) can be interpreted as the *market selection dynamics*. Strategies that have higher wealth than their competitors are considered to be fitter. If one strategy accumulates total wealth in the long term, it is said to be selected by the market. Since prices are a weighted combination of the strategies with weights equal to wealth, such a strategy asymptotically 'determines' asset prices: relative prices are asymptotically equal to the portfolio weights of a selected strategy.

The introduction of investment strategies, as well as the above definition of asset prices, marks a departure from the general equilibrium approach to market selection, where agents have demand functions and maximize utility over an infinite time-horizon (Blume and Easley, 1992 and 2006; Sandroni, 2000). Notwithstanding the apparent simplicity of the model, Amir et al. (2005) have shown that each equilibrium price path in such a general equilibrium model can be obtained by an appropriate specification of investment strategies whose portfolio weights are, in general, time- and history-dependent.

2.2 DYNAMICS OF INVESTMENT STRATEGIES

The wealth dynamics of the standard evolutionary finance model describe how the aggregate behavior changes over time for a fixed set of investment strategies. Here, the set of investment strategies is changing over time as well. This allows investigating whether the market mechanism is strong enough for the population to discover the Kelly rule, even though each individual investor lacks the analytical ability to do so.

The maximum number of investment strategies that can be active in the market at any period of time is limited to a finite number I . Each investment strategy $\lambda^i = (\lambda_1^i, \dots, \lambda_K^i)$ is represented by a program which is given in the form of a function

$$\tilde{\lambda}^i : \mathcal{S} \times \{1, \dots, K\} \rightarrow \mathbb{R}, \quad (2)$$

where $\mathcal{S}$ is the set of potential signals σ associated with the true state s . $\tilde{\lambda}_k^i(\sigma)$ is the non-normalized budget allocated to the purchase of asset k , given that the last observed signal is σ . Examples of programs are provided in Table I. In order to compute the portfolio weights for a given function $\tilde{\lambda}^i$, a normalization is carried out as follows. First define $\bar{\lambda}_k^i(\sigma) := \max\{0, \tilde{\lambda}_k^i(\sigma)\}$, and then let $\lambda_k^i(\sigma) := \bar{\lambda}_k^i(\sigma) / [\sum_{n=1}^K \bar{\lambda}_n^i(\sigma)]$. If the denominator is zero, we set $\lambda_k^i(\sigma) = 1$ for some randomly chosen k .

Two main cases are considered: (a) state-dependent strategies with complete information about states, i.e., $\mathcal{S} = \{1, \dots, S\}$, and (b) price-dependent strategies where the set of signals contains the last observed price vector corresponding to the present state, i.e., $\mathcal{S} = \{1, \dots, S\} \times \mathbb{R}^K$. In general the information set could be an indexed partition of the state space or contain other information such as asset payoffs and their moving averages.

2.3 IMPLEMENTATION BY GENETIC PROGRAMMING

Genetic programming (GP) (Koza, 1992; Smith, 1980) is a technique for programming computers by natural selection. It addresses the challenge of getting a computer to do what needs to be done without explaining everything

in detail, and it attempts to achieve this goal by breeding a population of programs using the principles of Darwinian selection and reproduction. The idea of mimicking evolution to simplify software design was introduced by Holland (1975), and it has been used successfully by computer scientists to solve problems in a variety of engineering fields. For our purpose, it is used as a model of decision making in the presence of tacit knowledge and bounded rationality.

GP-algorithms produce programs which consist of instructions to read or manipulate data. Each program produces some output, which the modeler interprets as the action taken by the program, and this action is evaluated to obtain a measure of the program's fitness. In our particular context, the action taken consists of the portfolio weights allocated to the K assets and the natural fitness measure is the accumulated wealth of the program.

In order to find the computer program which best solves a given task, the GP-algorithm starts by randomly generating a large population of programs. It then continues for a large number of iterations by replacing low performing programs with a genetic recombination of high performing ones. The standard genetic operators are crossover and mutation. These are explained below, along with an additional operator (noise) that is used here to test the robustness of our results.

GP-algorithms differ in the way they mimic Darwinian evolution. Here, a steady-state algorithm with tournament selection is used. It works as follows:

1. *Tournament*: Randomly select four programs from the pool and rank them according to their accumulated wealth.
2. *Reproduction*: Replace the two programs with lowest wealth with copies of the other two.
3. *Mutation*: Each of the two programs copied undergoes a mutation with probability μ : randomly select a single instruction from the program, and replace it with a randomly generated instruction.
4. *Crossover*: With probability χ , recombine the genetic material of the two copied (and possibly mutated) programs by swapping one randomly selected set of instructions from both programs.
5. *Noise*: Each of the two newly generated programs is replaced by a randomly selected program with probability η .

If a program has strictly positive wealth, it leaves the tournament with the same wealth. However, if a program has zero wealth, it is endowed with one percent of the average wealth which is given by $(1/I) \sum_{k=1}^K D^k(s_{t+1})$.

The algorithm is run for a population size of 1,000 and 250,000 iterations, with 20 tournament selections per iteration. The mutation and crossover probabilities are set to $\mu = 0.9$ and $\chi = 0.5$, respectively, and the noise probability η is varied from 0 to 0.96 in order to check the robustness of our results.

GP-algorithms also differ with respect to the type of building blocks they use to construct individual programs. One typically uses a subset of elementary instructions from some existing programming language, such as LISP, Java or machine code. In this paper, we use a machine code version of GP, which is introduced by Nordin (1997): each program consists of a list of machine instructions which operate on variables and constants stored in memory, using the CPU floating point registers to store and manipulate temporary variables.

Table I illustrates some aspects of the machine code GP algorithm by means of an example in which the maximum program length is limited to 6 instructions. In practice, the maximum program length is much larger; in our simulations, it consists of 128 instructions. The left part of the table depicts two programs, A and B, with 5 and 6 instructions, respectively. The right part shows the outcome of a crossover at instruction slots 3–6, which produces two new programs, C and D. R_0 and R_1 refer to floating point registers 0 and 1 of an Intel compatible CPU, which has a total of 8 such registers. The GP algorithm clears these registers by loading them with the value 0.0 before passing a program to the CPU for execution. Input variables consist of the state s and asset k (see Equation (2)), and the output is the content of register R_0 after all program instructions have been processed by the CPU. The output is interpreted as the non-normalized portfolio weight $\hat{\lambda}_k(s)$ for asset k in state s . For Program A, the *normalized* portfolio weights $\lambda_k(s)$ are constant and equal to $1/K$ for each state and asset, while for Program C they vary according to state and asset.

Instead of using GP to model the dynamics of investment strategies, a *genetic algorithm* (GA) (Holland, 1975) could be used in the first set of experiments. The main difference is that GA operates on vectors of numbers instead of vectors of program instructions. GA is applicable as long as the set of signals is finite (e.g., if only state and asset index are contained in the information set). Any investment strategy can then be represented as a vector of real numbers. Evolving such vectors with GA is simpler and computationally less demanding than evolving functions with GP. If strategies can also use past

Table I. Example of the program structure and the crossover operation

Instr. slot	Before crossover		After crossover	
	Program A	Program B	Program C	Program D
1	$R_0 = s$	$R_1 = k$	$R_0 = s$	$R_1 = k$
2	$R_1 = k$	$R_0 = R_0 - 2$	$R_1 = k$	$R_0 = R_0 - 2$
3	$R_1 = R_1 * R_0$	$R_0 = R_0 * R_0$	$R_0 = R_0 * R_0$	$R_1 = R_1 * R_0$
4	$R_0 = R_0 / 5$	$R_0 = R_0 / R_1$	$R_0 = R_0 / R_1$	$R_0 = R_0 / 5$
5	Return R_0	$R_0 = \min(R_0, R_1)$	$R_0 = \min(R_0, R_1)$	Return R_0
6		<i>Return R_0</i>	<i>Return R_0</i>	
$\hat{\lambda}_k(s)$	$s/5$	$\min(4/k, k)$	$\min(s^2/k, k)$	$-2/5$

prices, the information set is a continuum and so is the range of strategies. For these experiments, the full generality of GP, which provides us with functional relations, is needed to analyze the model.

3. The Experiments

Four sets of experiments are simulated to analyze the impact of a change in the market structure, as well as in the stochastic process that determines the state. The market is either complete (i.e., the rank of the payoff matrix is equal to the number of states) or incomplete (i.e., there are fewer assets than states). The state of nature is given either by an i.i.d. process or a Markov process.

In each experiment, there are three states of nature, $s = 1, 2, 3$, and $K = 3$ resp. $K = 2$ assets. The payoff matrix D is defined in Table II. Note that the complete market consists of Arrow securities.

In the i.i.d. case, all states have equal probability, i.e., $\pi_s = 1/3$ for $s = 1, 2, 3$. In the Markov case, however, the probability of the next period's state depends on the current state. The matrix of transition probabilities is given by

$$\Pi = \begin{pmatrix} 0.7 & 0.2 & 0.1 \\ 0.1 & 0.7 & 0.2 \\ 0.2 & 0.1 & 0.7 \end{pmatrix}. \quad (3)$$

The stationary distribution of this Markov process, denoted by ρ , is given by $\rho_1 = \rho_2 = \rho_3 = 1/3$.

The Kelly rule in each of these cases is given by the expected values of the assets' relative payoffs:

$$\lambda_k^*(s) = \sum_{u=1}^S \Pi_{su} R^k(u), \quad \text{where } R^k(u) = \frac{D^k(u)}{\sum_{n=1}^K D^n(u)}, \quad (4)$$

with $k = 1, \dots, K$ and $s = 1, \dots, S$. In the i.i.d. case, the current state does not impact the probability of the next period's state, i.e., $\Pi_{su} \equiv \pi_u$, and, therefore, the Kelly rule is a constant vector. Table III summarizes the numerical values of the Kelly rule in the four experiments.

The investment strategy λ^* provides the equilibrium prediction for the long-run outcome of individual behavior, as well as for the asymptotic values

Table II. Payoff matrix D in the complete, resp. incomplete, market case

Complete market	Incomplete market
$D = \begin{pmatrix} 1 & 0 & 0 \\ 0 & 2 & 0 \\ 0 & 0 & 3 \end{pmatrix}$	$D = \begin{pmatrix} 1 & 1 \\ 1 & 1 \\ 0 & 3 \end{pmatrix}$

Table III. Kelly rule in the four experiments

IID		Markov	
Complete	Incomplete	Complete	Incomplete
$\lambda^* \equiv \left(\frac{1}{3}, \frac{1}{3}, \frac{1}{3}\right)$	$\lambda^* \equiv \left(\frac{1}{3}, \frac{2}{3}\right)$	$\lambda^*(s=1) = (.7, .2, .1)$ $\lambda^*(s=2) = (.1, .7, .2)$ $\lambda^*(s=3) = (.2, .1, .7)$	$\lambda^*(s=1) = (.45, .55)$ $\lambda^*(s=2) = (.40, .60)$ $\lambda^*(s=3) = (.15, .85)$

of (relative) asset prices. The two main questions to be investigated are (1) whether the competitive process of genetic programming is powerful enough to drive the population towards the Kelly rule and, if yes, (2) whether this convergence is robust against noise.

To answer the first question, the distance between the Kelly rule and the wealthiest investment strategy is measured, as well as the distance to the market prices. Prices are equal to the average strategy of the population because they correspond to the wealth-weighted average, as explained in Section 2.1. The wealthiest strategy is defined here as the sequence of investment strategies that is provided by selecting, in any one period in time, the strategy with the highest wealth. Two measures are applied: (a) the Euclidean distance and (b) the expected growth rate. The latter quantity measures the growth potential of an investment strategy relative to a benchmark. Two benchmarks are employed: the Kelly rule and the current state of the market, which is given by the current price system.

The second question on the robustness of the convergence is analyzed by comparing long-term outcomes for different values of the average number of randomly generated programs. This is achieved by varying the noise probability η , see Step 5 of the GP-algorithm.

3.1 SIMULATION RESULTS

Each experiment consists of a simulation of a population of $I = 1,000$ genetic programs. Programs are initialized by randomly chosen functions, which are arrays of random length, filled with random draws from the instruction set. The noise probability η is varied, in each experiment, between 0 and 0.96 with increments of 0.04. The Darwinian finance model is run for a total of 250,000 periods for each set of parameters.

The distance measures are defined as follows. The probability-weighted Euclidean distance between a strategy λ^i and the Kelly rule λ^* is given by

$$d^*(\lambda^i) := \sum_{s=1}^S \rho_s \sqrt{\sum_{k=1}^K (\lambda_k^i(s) - \lambda_k^*(s))^2}. \quad (5)$$

In the i.i.d. case, the stationary distribution ρ is given by $\rho_s = \pi_s$.

The second distance measure can be defined in terms of the growth rate of a strategy λ^i relative to the Kelly benchmark, $g^*(\lambda^i)$, or its growth rate relative to the market, $g^M(\lambda^i)$. The first one is defined as

$$g^*(\lambda^i) := \exp \left(\sum_{s=1}^S \rho_s \sum_{u=1}^S \Pi_{su} \ln \left(\sum_{k=1}^K \frac{R^k(u)}{\lambda_k^*(s)} \lambda_k^i(s) \right) \right), \quad (6)$$

where the relative payoffs $R^k(u)$ are given by (4) and ρ is the stationary distribution of the stochastic process that determines the state of nature. In the i.i.d. case, one has $\Pi_{su} \equiv \pi_u$ and $\lambda_k^*(s) \equiv \lambda_k^*$ in Equation (6). Taking the exponential allows a comparison of growth rates in terms of percentages. The standard definition is obtained by dropping the exponential function (see Hens and Schenk-Hoppé, 2005). One could also try to measure the distance by the relative entropy of strategy λ^i and the Kelly rule. In the case of Arrow securities the relative entropy is equal to the expected logarithmic growth rate of λ^i . However, if the market is incomplete, the relative entropy, in general, provides no information about a strategy's growth rate (see, e.g., Blume and Easley, 2006; Sandroni, 2005). The growth rate relative to the market, $g^M(\lambda^i)$, is defined by replacing $\lambda_k^*(s)$ in (6) by the (moving) benchmark $\lambda_k^M(s) := p_k(s) = q_k(s) / [\sum_{n=1}^K q_n(s)]$. $p_k(s)$ is the relative price of asset k .

The simulation results for these different measures are summarized in Tables IV and V. In both IID experiments, investment behavior converges quickly to the Kelly benchmark for all noise levels. This can be seen from

Table IV. Summary statistics for 4 sets of experiments without price information

The table contains sample statistics across all runs and iterations for relative growth rates and distances from the Kelly benchmark. λ^W denotes the wealthiest strategy and λ^M denotes the market portfolio. The number of observations for each variable is 6,250 (250 samples of 25 runs).

Variable	IID				Markov			
	Mean	Std.dev	Min	Max	Mean	Std.dev	Min	Max
Complete market								
$d^*(\lambda^W)$	0.001	0.022	0.000	0.817	0.054	0.048	0.003	1.128
$d^*(\lambda^M)$	0.001	0.006	0.000	0.305	0.053	0.044	0.002	0.715
$g^*(\lambda^W)$	0.999	0.038	0.000	1.000	0.988	0.033	0.000	1.000
$g^M(\lambda^W)$	0.998	0.044	0.000	1.005	0.996	0.058	0.000	1.008
Incomplete market								
$d^*(\lambda^W)$	0.003	0.009	0.000	0.196	0.061	0.032	0.014	0.374
$d^*(\lambda^M)$	0.003	0.007	0.000	0.179	0.023	0.017	0.001	0.251
$g^*(\lambda^W)$	1.000	0.000	0.990	1.000	0.999	0.001	0.959	1.000
$g^M(\lambda^W)$	0.999	0.031	0.000	1.009	0.989	0.098	0.000	1.010

Table V. Average wealth of 5 groups of investment strategies in the 4 sets of experiments without price information

This table shows the proportion of wealth owned by the wealthiest strategy (λ^W) and by those strategies that are close to the Kelly rule in terms of relative growth rates ($g^* \geq . . .$). Numbers are averages taken across 6,250 observations for each experiment.

Dividend process	Market structure	λ^W	$g^* \geq 0.950$	$g^* \geq 0.975$	$g^* \geq 0.990$	$g^* \geq 0.999$
IID	Complete	0.040	0.999	0.999	0.998	0.998
	Incomplete	0.789	0.997	0.997	0.995	0.991
Markov	Complete	0.568	0.981	0.926	0.538	0.217
	Incomplete	0.528	0.996	0.994	0.988	0.625

the small means and standard deviations of the two distance measures $d^*(\lambda^W)$ and $d^*(\lambda^M)$ in Table IV. In the Markov experiments, the mean distances from the Kelly benchmark are higher. Nonetheless, the wealthiest strategy performs quite well relative to the market, as well as to the Kelly benchmark, as manifested by the high relative growth rates $g^*(\lambda^W)$ and $g^M(\lambda^W)$. Table V shows, however, that the wealthiest strategy does not possess the entire market wealth. A substantial amount of wealth is managed by other investment strategies (for instance nearly 50% in both Markov experiments), but those strategies are also quite close to the Kelly rule in terms of relative growth rates, as can be seen from the 4 rightmost columns of Table V.

In the Markov cases, particularly in the incomplete market setting, the price system is closer to the Kelly benchmark than the wealthiest investment strategy, i.e., the market is “smarter” than the most successful investor, cf. Table IV. This finding is in line with Bossaerts et al. (2005). In experiments with human subjects, they find that, while asset prices converge quickly and agree—by and large—with those of the CAPM, the portfolio choice predictions of this theory remain significantly off target throughout the experiment.

We analyze the Markov experiments in more detail by regressing distances on noise and simulation time in a panel data model with random effects and autoregressive errors. The estimated relationship is

$$\ln(d_{rt}^*) = \beta_0 + \beta_1 \ln(t/250) + \beta_2 \ln(1 - \eta_r) + u_r + \epsilon_{rt}, \tag{7}$$

where u_r is a run specific error term and ϵ_{rt} is AR(1). d_{rt}^* is the Euclidean distance from the Kelly rule at iteration $1000 \cdot t$ of run r and η_r is the noise level in run r . Table VI contains the results. Estimated parameters are given in Panel (A) and predicted distances for selected values of the explanatory variables are given in Panel (B).

The negative coefficients of $\ln(t/250)$ and $\ln(1 - \eta_r)$ in panel (A) show that distances from the Kelly rule decrease over time and that the process slows down when the noise level is increased. Panel (B) gives the predicted

Table VI. Model and predicted values of Euclidean distance from the Kelly rule. Markov dividend process.

The estimated model is $\ln(d_{rt}^*) = \beta_0 + \beta_1 \ln(t/250) + \beta_2 \ln(1 - \eta_r) + u_r + \epsilon_{rt}$, where u_r is a run specific error term and ϵ_{rt} is an AR(1) process. d_{rt}^* is the Euclidean distance from the Kelly rule at iteration $1000 \cdot t$ of run r and η_r is the run r noise level. λ^W denotes the wealthiest strategy and λ^M denotes the market portfolio. Standard errors are given in parentheses. * and ** denote significantly different from zero at the 5% and 1% level, respectively (two-tailed test).

Dependent variable	(A) Parameters				(B) Predicted distance			
	β_0	β_1	β_2	R^2	$t \backslash \eta$	0.9	0.5	0.0
Complete market								
$d^*(\lambda^W)$	-4.238**	-0.585**	-0.490**	0.474	5	0.439	0.200	0.142
	(0.139)	(0.019)	(0.113)		250	0.045	0.020	0.014
$d^*(\lambda^M)$	-4.220**	-0.541**	-0.502**	0.477	5	0.387	0.172	0.122
	(0.145)	(0.018)	(0.119)		250	0.047	0.021	0.015
Incomplete market								
$d^*(\lambda^W)$	-3.246**	-0.170**	-0.190*	0.196	5	0.117	0.087	0.076
	(0.109)	(0.006)	(0.090)		250	0.060	0.044	0.039
$d^*(\lambda^M)$	-4.515**	-0.254**	-0.323**	0.284	5	0.062	0.037	0.030
	(0.073)	(0.008)	(0.059)		250	0.023	0.014	0.011

(shrinking) distance to the Kelly rule for noise levels 0, 0.5 and 0.9. The most striking observation is that, in all cases, the wealthiest investment strategy, as well as the market, converges to the Kelly rule. Even for high values of the noise parameter, convergence prevails, though at a lower speed.

Panel (B) of Table VI also shows that the market is closer to the Kelly rule than the wealthiest investment strategy, in particular at the beginning of each experiment. The effect is particularly pronounced in the incomplete market case with Markovian states. This observation implies that there are strategies in the market that move prices closer to the Kelly rule, though they do not provide the wealthiest strategy in the market.

3.2 COMPLETE MARKET WITH MARKOV STATES

This section provides a more detailed analysis of the case with a complete market and a Markov payoff process. The goal of this exercise is to obtain a better insight into the mechanisms that drive the Darwinian dynamics which give rise to the findings reported above. This particular case is chosen because it exhibits the highest deviation of prices from the Kelly rule. The noise probability is set to 20% throughout the following.




Figure 1. Euclidean distance between the market prices resp. the wealthiest strategy and the Kelly rule.

Figure 1 shows the distance between the Kelly benchmark and the market, $d^*(\lambda^M)$, as well as the wealthiest strategy, $d^*(\lambda^W)$, for one run of the simulation, see (5). Both distance measures decrease almost monotonically, and the descent mainly follows a step function. In this leapfrogging movement towards the benchmark, the wealthiest strategy is overtaken, every now and again, by a better performing competitor. This occurs only nine times during the time horizon considered in Figure 1. This observation implies that, most of the time, the wealthiest strategy is not the rule with the highest growth rate. The competing strategy that eventually replaces the wealthiest rule is typically much closer to the Kelly rule, as shown by the size of the steps in the convergence. If the wealthiest strategy is close to the Kelly rule, this process takes a considerable amount of time. The reason is that genetic programming creates better, but somewhat poorer, strategies from the genetic material of the wealthiest rule rather than improving the currently richest strategies. After 250,000 iterations, the wealthiest strategy is very similar to the Kelly rule, with a distance of only 0.003. This number is in agreement with the results reported in Table VI.

The distance between market prices and the Kelly rule displays similar behavior though, in most periods, this measure is closer to the Kelly rule than the wealthiest strategy. Only in the last 20,000 iterations is the market price further away from the benchmark than the wealthiest strategy. This particular effect is caused by the relatively small perturbation through the noise component of the tournament, which has a stronger impact when the population is very close to the benchmark.

On the aggregate level, interest focuses on the convergence dynamics of asset prices. Figure 2 shows the relative price of Asset 1 for each period of time and for each state $s = 1, 2, 3$ (one of which is revealed in any given period) for a




Figure 2. Relative price of Asset 1 by state and iteration. Horizontal lines correspond to the Kelly benchmark.

simulation of one run of 250,000 iterations. The benchmark, which is given by the Kelly rule, is the set of relative asset prices $p_1^*(1) = 0.7$, $p_1^*(2) = 0.2$ and $p_1^*(3) = 0.1$ in states 1, 2 and 3, respectively. The simulated prices are some distance from the benchmark in the first 150,000 iterations, but they exhibit a clear tendency to converge to this benchmark. After approximately 175,000 iterations, the relative prices are almost identical to those derived from the Kelly rule. However, systematic mispricing occurs over long time horizons. For instance, in state 1, the asset is first undervalued and later overvalued. Moreover, the prices in all states can simultaneously be too low compared to the benchmark. In Figure 2, this occurs during iterations 30,000–40,000. Relative prices of an asset do not necessarily sum up to one across states, but of course, the sum across relative asset prices is equal to one in each state.

A proxy for the genetic material that is present in the population is provided by the distribution of growth rates relative to the Kelly rule. These values are, in fact, only potential growth rates because many strategies have zero wealth. Figure 3 reports the distribution of the individual growth rates relative to the Kelly rule for the entire population across iterations. An analysis of the distribution of growth rates relative to the current market prices gives a similar picture to the one provided in Figure 3.

Programs with a zero relative growth rate (defined in (6)) consist of strategies that are not fully diversified; a behavior that carries a positive probability of losing all one's wealth. The second group consists of strategies with positive relative growth rates below 0.95. A large proportion of these strategies have constant and equal portfolio weights, a strategy which yields a relative growth rate of 0.7432. These programs simply return the same constant for each asset and state. They are quite successful in the first 1,000 iterations of the simulation, because they avoid mistakes, such as under-diversification




Figure 3. Distribution of growth rates in the population relative to the Kelly rule. (Data are exponentially smoothed, with a parameter value of 0.2.).

or extreme deviations from actual probabilities. Since the probability of generating such a strategy from scratch or by crossover is high, the population contains a large proportion of those in the early stages of the simulation.

The distribution of relative growth rates changes significantly over time. The number of strategies with equal weights is quickly reduced from an initially high level. The reason is that, as the fraction of efficient strategies increases, the equal weight strategies lose money faster, which reduces their chances of reproducing. Halfway into the simulation, strategies with relative growth rates above 0.999 begin to take over and, after 250,000 iterations, they make up about 40% of the population. This illustrates the ability of the GP-algorithm to generate extremely efficient investment strategies.

Interestingly, the fraction of strategies with a zero relative growth rate increases from approximately 20% to 30% during the simulation. Since these strategies are not fully diversified, they act like gamblers who take extreme positions in some assets. We believe this result is driven by the following evolutionary mechanism: since wealth is a prerequisite for reproduction, the struggle for survival is basically a struggle to get rich. As the population becomes dominated by extremely efficient strategies, it grows harder to get rich by investing prudently. This makes gambling for large stakes increasingly attractive as it can open a window of opportunity to reproduce. By this mechanism, competition produces not only more winners but also more losers, which is in agreement with Figure 3.

4. Price-dependence and Market Dynamics

This section presents a generalization of the previous setting: strategies are given access to information on past prices. The main issue is how and to what

extent genetic programs use this additional information. To obtain robust results, a case is analyzed in which the growth-optimal portfolio coincides with the Kelly rule for all prices, but where the latter is dependent on the state of nature. Prices should not matter for a growth-oriented investor, though states should. We focus on the experiment with a complete market and a Markov dividend process, studied in detail for state-dependent strategies in Section 3.2.

The availability of price information indeed gives rise to new types of behavior. The simulation results reported below show in particular that a prominent role is played by the market portfolio. Its importance is due to the fact that buying the market portfolio entails a constant share of total wealth and hence increases its probability of survival and reproduction relative to a strategy with more volatile payoffs.

In order to see that buying the market portfolio preserves relative wealth, consider a strategy, denoted λ^M , whose portfolio weights correspond to the relative prices $p_{k,t}$ (which are the market portfolio weights):

$$\lambda_{k,t}^M = p_{k,t} := \frac{q_{k,t}}{\sum_{n=1}^K q_{n,t}}. \quad (8)$$

Let w_t^M denote the wealth of strategy λ^M . The portfolio holdings of this strategy are equal to

$$\theta_{k,t}^M := \frac{\lambda_{k,t}^M w_t^M}{q_{k,t}} = \frac{w_t^M}{\sum_{n=1}^K q_{n,t}} = \frac{w_t^M}{W_t} = r_t^M,$$

where $W_t := \sum_{i=1}^I w_t^i = \sum_{n=1}^K q_{n,t}$ is the total wealth and r_t^M is the relative wealth of strategy λ^M . It holds the same amount of units of each asset because $\theta_{k,t}^M$ is independent of k . Its payoff can be computed as

$$w_{t+1}^M = \sum_{k=1}^K D^k(s_{t+1}) \theta_{k,t}^M = \sum_{k=1}^K D^k(s_{t+1}) r_t^M = \bar{D}(s_{t+1}) r_t^M,$$

with $\bar{D}(s_{t+1})$ denoting the aggregate dividend payoff in state s_{t+1} . It follows from (1) that $\bar{D}(s_{t+1}) = W_{t+1}$. Thus, an investor who buys the market portfolio has constant relative wealth, i.e., $r_{t+1}^M = r_t^M$.

To introduce price dependency, an investment strategy is defined here as a function $\tilde{\lambda} : \{1, \dots, S\} \times \mathbb{R}^K \times \{1, \dots, K\} \rightarrow \mathbb{R}$, where $\tilde{\lambda}_k(s, \hat{p}(s))$ is the (non-normalized) portfolio weight of asset k given that the current state is s . The information set is $\mathcal{S} = \{1, \dots, S\} \times \mathbb{R}^K$. The new argument $\hat{p}(s)$ denotes the most recently observed relative price vector corresponding to the current state s . For instance, if $(s_{t-3}, s_{t-2}, s_{t-1}, s_t) = (1, 1, 2, 1)$, where $s = s_t = 1$, then $\hat{p}(s)$ is the vector of relative prices observed in period $t-2$. Trading is subject to slippage because strategies only have access to the last observed price vector (i.e., for the current state) rather than the current market clearing prices. This

slippage comes from the time elapsing from order placement to execution, which precludes the purchase of a perfect market portfolio. However, if prices exhibit low volatility, the strategy defined in (8) delivers an asset allocation that corresponds well with the market portfolio. In periods of high volatility, the fit is not so good. Despite the slippage, this proxy of the market portfolio provides an opportunity to protect one's investment from downside risk at the cost of missing out on its upside. Carrying out a market portfolio investment is actually a simple task for the genetic programs: they only have to output the price of the relevant assets, whereas any other behavior requires computation. The following shows that successful investment strategies make skillful use of this investment opportunity.

Extending the information set from states to states and prices more than doubles the number of input variables to any given strategy. The new data are vectors of real numbers rather than elements of a finite set. This increases the complexity faced by the GP-algorithm considerably. Since this is likely to slow down convergence, the effect of doubling the population size to 2,000 is tested. Table VII presents the result of applying the regression model in Equation (7) to data from a simulation with 25 runs, with noise levels between 0 and 96%, and population sizes of 1,000 and 2,000. The increase in population size indeed compensates for the increase in complexity, cf. Table VI. As in the corresponding base case considered in Section 3.2, the noise probability is set to 20%.

Table VII. Model and predicted values of Euclidean distance from the Kelly rule for price-dependent strategies. Markov dividend process and complete market.

The estimated model is $\ln(d_{rt}^*) = \beta_0 + \beta_1 \ln(t/250) + \beta_2 \ln(1 - \eta_r) + u_r + \epsilon_{rt}$, where u_r is a run specific error term and ϵ_{rt} is an AR(1) process. d_{rt}^* is the Euclidean distance from the Kelly rule at iteration $1000 \cdot t$ of run r and η_r is the run r noise level. λ^W is the wealthiest strategy and λ^M is the market portfolio. Standard errors are given in parentheses. * and ** denote significantly different from zero at the 5% and 1% level, respectively (two-tailed test).

Dependent variable	(A) Parameters				(B) Predicted distance			
	β_0	β_1	β_2	R^2	$t \setminus \eta$	0.9	0.5	0.0
Population size of 1,000								
$d^*(\lambda^W)$	-3.354**	-0.340**	-0.239**	0.259	5	0.229	0.156	0.132
	(0.074)	(0.011)	(0.060)		250	0.061	0.041	0.035
$d^*(\lambda^M)$	-3.360**	-0.336**	-0.241**	0.253	5	0.226	0.153	0.129
	(0.074)	(0.011)	(0.060)		250	0.061	0.041	0.035
Population size of 2,000								
$d^*(\lambda^W)$	-3.911**	-0.482**	-0.334**	0.392	5	0.285	0.166	0.132
	(0.067)	(0.013)	(0.054)		250	0.043	0.025	0.020
$d^*(\lambda^M)$	-3.934**	-0.484**	-0.338**	0.388	5	0.283	0.164	0.130
	(0.068)	(0.013)	(0.055)		250	0.043	0.025	0.020

The dynamics of the Euclidean distance between the wealthiest strategy and the Kelly rule, reported in Figure 4, are strikingly different from the base case depicted in Figure 1. Volatility is much higher, there is no leap-frogging effect and, despite the variability of prices, the market is not closer to the Kelly rule than the wealthiest rule. Moreover, there are occasional bursts of large deviations from a generally volatile trend toward smaller distances, cf. Table VII.

The dynamics of prices are captured in Figure 5, which depicts one time series of 250,000 iterations' length for the relative price of Asset 1 for each of the states $s = 1, 2, 3$. The benchmark is given by the Kelly prices $p_1^*(1) = 0.7$,




Figure 4. Euclidean distance between the Kelly rule and the wealthiest strategy resp. the market prices (which are virtually identical here) for price-dependent strategies.




Figure 5. Relative price of Asset 1 by state and iteration (in 1,000s) for price-dependent strategies and corresponding Kelly benchmark (horizontal lines).

$p_1^*(2) = 0.2$ and $p_1^*(3) = 0.1$ in the respective states, cf. Table III. Figure 5 shows that prices have a clear tendency to vary around their respective Kelly benchmark levels, except for a short initial period of 2,000 iterations, during which significant mispricing occurs. A comparison with the corresponding case without price information (Figure 2) reveals that the availability of price information creates a considerable amount of persistent price volatility.

In order to obtain an understanding of the market dynamics with price-dependent strategies, a careful analysis of the successful investment strategies and their market shares is necessary. It turns out that good performing strategies use price information in a sophisticated way while, at the same time, being composed of similar “building blocks.” In particular, usage of the market portfolio is wide-spread. Averaging across all 250,000 iterations, 30% of the population (representing about 32% of the total wealth) buy the market portfolio. The strategies which use *some* price information, without necessarily buying the market portfolio, amount to 97.5% of the population and represent about 99.99% of the total wealth on average. In the setting with price-dependent strategies, it seems that price information is necessary for survival.

Individual strategies can be classified according to their behavioral variation over time. Some rules follow a *passive* strategy and invest 100% of their wealth in the market portfolio at all times. Others adopt an *active* strategy by always deviating from the market portfolio. The third group of rules applies a *hybrid* strategy: in some periods, they passively buy the market portfolio, while in others they take some active risk by deviating from the market benchmark. Figure 6 applies this classification to the *observed* behavior of investment strategies. It shows how the proportions of wealth managed by each of these three sets of rules varies over time. Note that active and passive strategies may have the potential to change their behavior, even if they did not do so in the




Figure 6. Percentage of wealth managed by different types of price-dependent strategies.

past, because prices did not take on those values which might trigger their latent behavior. The classification therefore gives upper limits on the number of active and passive strategies and provides a lower limit for the number of hybrid strategies. Figure 6 also supplies information on the market share of strategies that currently take active positions. This set of strategies consists of all of the active rules and some of the hybrid rules.

There is considerable variation in the number of hybrid strategies deviating from the market portfolio investment. In the early rounds, active strategies gain and passive strategies lose in terms of market shares. After approximately 100,000 iterations, there is a steep increase in the market share of hybrid strategies. This occurs because a few rules, which were previously classified as active, bought the market portfolio for the first time during a short period of severe mispricing. This is mainly due to price dynamics because, as discussed in detail below, strategies typically have a "trigger," which induces a switch from passive to active investment behavior if prices enter a certain region of the price space.

Other large jumps in the distribution of strategy types (at 130,000, 165,000 and 200,000) are preceded by periods of unusually high volatility, and large deviations from the Kelly rule. This typically occurs when some extremely risky (e.g., under-diversified) strategy experiences a streak of good luck and becomes wealthy enough to have some market impact before it eventually goes bankrupt. For normal levels of price volatility, the presence of passive strategies amplifies its market impact and contributes to the volatility. For high levels of volatility, information about past prices gives imprecise information about the current market portfolio. This increases the risk for those strategies that rely heavily on price information. The gains and losses incurred may change the wealth distribution, alter the aggregate behavior and cause relative prices to settle down at new levels. The sharp decline around period 165,000 in the market share of currently active strategies is triggered by a change in relative price levels, following a volatility shock of the type just described. Figure 6 indicates that the main part of this decline can be attributed to hybrid strategies, which change their mode of behavior from active to passive.

In order to investigate the mechanism which causes this shift in behavior, we provide a detailed analysis of one specific strategy from the experiment—the wealthiest strategy at the last iteration of the experiment. This strategy, denoted $\lambda^{LW}(s, p)$, is present in the population for more than the last 150,000 iterations. Because the portfolio weights of λ^{LW} are a function of the state s and the vector of relative prices $p = (p_1, p_2, p_3)$, the rule is therefore of the hybrid type. In states 1 and 2, behavior remains passive throughout, but in state 3, it switches between active and passive modes. Figure 7 summarizes its behavior in state 3. In the remainder of this section, we fix $s = 3$ and refer to $\lambda^{LW}(s, p)$ and $\lambda^*(s)$ as $\lambda^{LW}(p)$ and λ^* , respectively.

Figure 7 is a contour plot of the distance between the market portfolio and the portfolio weights of strategy λ^{LW} in state 3. Market portfolios, which




Figure 7. Contour plot of $\delta(p) := \|\lambda^{LW}(p) - p\|$ for relative prices $p = (p_1, p_2, p_3)$ in the unit simplex. $\delta(p)$ is the Euclidean distance between the market portfolio p and the portfolio weights of strategy λ^{LW} in state 3. Darker areas represent larger values of $\delta(p)$, and white areas represent those prices for which λ^{LW} is in perfect agreement with the market portfolio, i.e., $\delta(p) = 0$. Kelly prices and Kelly portfolio weights are given by $\lambda^* = (0.2, 0.1, 0.7)$.

coincide with vectors of relative prices, are represented as points in the unit simplex. The point $\lambda^* = (0.2, 0.1, 0.7)$ is the Kelly benchmark portfolio in state 3 (see Table III). For each price vector $p = (p_1, p_2, p_3)$, $\lambda^{LW}(p)$ are the portfolio weights of strategy λ^{LW} , and $\delta(p) := \|\lambda^{LW}(p) - p\|$ is the Euclidean distance between these two. In the figure, darker shadings corresponds with higher values of $\delta(p)$.

Strategy λ^{LW} is in the passive mode (P) in a subset of the simplex that consists of two parts: (1) the line segment defined by all price vectors p with $p_2 = \lambda_2^* = 0.1$ and (2) the white triangular area along the north-east border of the simplex, which corresponds with low prices for Asset 1. The remaining area of the simplex constitutes the active mode (A). The switch between active and passive modes is triggered by low prices of Asset 1. The functional form of strategy λ^{LW} in the two modes is given by

$$\lambda^{LW}(p) = \begin{cases} (1 - \lambda_2^* - \epsilon(p_2)) \left(\frac{p_1}{p_1 + p_3}, \frac{\lambda_2^* + \epsilon(p_2)}{1 - \lambda_2^* - \epsilon(p_2)}, \frac{p_3}{p_1 + p_3} \right) & \text{if } p \in A; \\ (p_1, p_2, p_3) & \text{if } p \in P. \end{cases} \quad (9)$$

The function $\epsilon(p_2)$ is convex with values $\epsilon(0) = 0.008$, $\epsilon(\lambda_2^*) = 0.000$ and $\epsilon(1) = -0.016$.

In the active mode, strategy λ^{LW} chooses a portfolio weight of Asset 2 that is close to the Kelly benchmark. For $p_2 \neq \lambda_2^*$, function $\epsilon(\cdot)$ provides an enhancement which amounts to small long or short positions in Asset 2 according to whether this asset is cheap or expensive relative to its Kelly

benchmark. The deviation from the Kelly weight λ_2^* is of the magnitude 10%. The remaining wealth is invested in the other two assets according to their price ratio, i.e., in the corresponding market portfolio of Assets 1 and 3. Strategy λ^{LW} is discontinuous on the border between active and passive modes (which marks the switch to passive investment for low prices of Asset 1), except along the line segment where $p_2 = \lambda_2^*$. However, Equation (9) reveals that the discontinuities match up to maintain equal ratios of portfolio weights and market prices for Assets 1 and 3.

The long-lived strategy λ^{LW} is specialized in the sense that it only takes active positions in one particular asset and in one particular state; otherwise it is passive and buys the market portfolio. A careful analysis of other wealthy investment strategies shows that most of them have exactly the same structure. They exploit mispricing in a nearly identical fashion to the strategy studied in detail above, though for a different asset and a different state. In other situations, they are passive and buy the market portfolio.

In our Darwinian model, the tournament process punishes investors with low wealth by a high probability of deletion. Avoiding severe losses is therefore important for survival, a goal that can be achieved by investing part of one's wealth in the market portfolio. Moreover, the logarithm of strategy λ^{LW} 's return has a smaller variance than the Kelly rule for about 80% of the time in the runs reported in Figures 5 and 6. The genetic recombination process favors this type of prudent investment behavior because extinction is more likely for rules that have a high volatility of returns.

Specializing in active investment under particular circumstances, which would also make sense from a practical point of view, stems from the genetic process that distributes trading skills across the population. The basic building blocks, which embody particular specialization, spread throughout the population of behavior rules. It is extremely unlikely, however, to have them all combined in one single rule, due to the random nature of genetic recombination in the tournament process and the lack of individual learning. Though complete knowledge is embodied in the population, this precludes the birth of a 'super trader.' Even though successful investors are more cautious when strategies are price- rather than just state-dependent, convergence to the Kelly rule prevails because of the push from the population's aggregate behavior towards the equilibrium prediction.

5. Conclusion

In this paper, we studied an evolutionary model of a financial market with short-lived assets, which comprises the two Darwinian processes of selection and reproduction. The framework generalizes an evolutionary finance model by applying a genetic programming approach to the creation of variety and the dynamics of investment behavior. The well-known Kelly rule serves as the

benchmark for optimal investment and asset prices. Our findings fully confirm the predictions derived from the analytic results in models that neglect the reproductive process which generates diverse behavior. In all experiments, the long-run outcome shows the emergence of investment styles and market prices in line with the Kelly rule. When price information is available in addition to information about the current state, strategies make sophisticated use of the data by employing the price information to invest part of their wealth in the market portfolio. This wide-spread behavior provides protection against severe losses and increases the chance of survival of the corresponding genetic material. The combination of this type of investment behavior with small bets on the Kelly rule eventually drives market prices to the Kelly benchmark.

References

- Algoet, P. H. and Cover, T. M. (1988) Asymptotic optimality and asymptotic equipartition properties of log-optimum investment, *Annals of Probability* **16**, 876–898.
- Allen, F. and Karjalainen, R. (1999) Using genetic algorithms to find technical trading rules, *Journal of Financial Economics* **51**, 245–271.
- Amir, R., Evstigneev, I. V., Hens, T., and Schenk-Hoppé, K. R. (2005) Market selection and survival of investment strategies, *Journal of Mathematical Economics* **41**, 105–122.
- Arifovic, J. (1996) The behavior of the exchange rate in the genetic algorithm and experimental economies, *Journal of Political Economy* **104**, 510–541.
- Blume, L. and Easley, D. (1992) Evolution and market behavior, *Journal of Economic Theory* **58**, 9–40.
- Blume, L. and Easley, D. (2006) If you are so smart, why aren't you rich? Belief selection in complete and incomplete markets, *Econometrica* **74**, 929–966.
- Bossaerts, P., Plott, C., and Zame, W. R. (2005) Prices and portfolio choices in financial markets: Theory and experiment, unpublished, UCLA Economics working paper No. 840, University of California.
- De Long, J. B., Shleifer, A., Summers, L. H., and Waldmann, R. J. (1990) Noise trader risk in financial markets, *Journal of Political Economy* **98**, 703–738.
- De Long, J. B., Shleifer, A., Summers, L. H., and Waldmann, R. J. (1991) The survival of noise traders in financial markets, *Journal of Business* **64**, 1–19.
- Evstigneev, I. V., Hens, T., and Schenk-Hoppé, K. R. (2002) Market selection of financial trading strategies: Global stability, *Mathematical Finance* **12**, 329–339.
- Evstigneev, I. V., Hens, T., and Schenk-Hoppé, K. R. (2006) Evolutionary stable stock markets, *Economic Theory* **27**, 449–468.
- Friedman, M. (1953) *Essays in Positive Economics*, The University of Chicago Press, Chicago.
- Hens, T. and Schenk-Hoppé, K. R. (2005) Evolutionary stability of portfolio rules in incomplete markets, *Journal of Mathematical Economics* **41**, 43–66.
- Holland, J. (1975) *Adaptation in Natural and Artificial Systems*, The University of Michigan Press, Ann Arbor.
- Hommes, C. (2001) Financial markets as nonlinear adaptive evolutionary systems, *Quantitative Finance* **1**, 149–167.
- Kelly, J. L. (1956) A new interpretation of information rate, *Bell System Technical Journal* **35**, 917–926.
- Kogan, L., Ross, S., Wang, J., and Westerfield, M. (2006) The survival and price impact of irrational traders, *Journal of Finance* **61**, 195–229.

- Koza, J. (1992) *Genetic Programming. On the Programming of Computers by Means of Natural Selection*, The MIT Press, Cambridge, Massachusetts.
- Lensberg, T. (1999) Investment behavior under Knightian uncertainty—An evolutionary approach, *Journal of Economic Dynamics and Control* **23**, 1587–1604.
- Neely, C. J., Weller, P. A., and Dittmar, R. (1997) Is technical analysis in the foreign exchange market profitable? A genetic programming approach, *Journal of Financial and Quantitative Analysis* **32**, 405–426.
- Nordin, P. (1997) *Evolutionary Program Induction of Binary Machine Code and its Applications*, Krehl Verlag, Münster.
- Polanyi, M. (1967) *The Tacit Dimension*, Routledge, London.
- Sandroni, A. (2000) Do markets favor agents able to make accurate predictions?, *Econometrica* **68**, 1303–1341.
- Sandroni, A. (2005) Market selection when markets are incomplete, *Journal of Mathematical Economics* **41**, 91–104.
- Shapley, L. and Shubik, M. (1977) Trade using one commodity as a means of payment, *Journal of Political Economy* **85**, 937–968.
- Smith, S. F. (1980) A learning system based on genetic adaptive algorithms, PhD Thesis, Computer Science Department, University of Pittsburgh.

Copyright of Review of Finance is the property of Oxford University Press - Journals Department and its content may not be copied or emailed to multiple sites or posted to a listserv without the copyright holder's express written permission. However, users may print, download, or email articles for individual use.