Journal of Economic Dynamics & Control
23 (1999) 1587–1604

JOURNAL OF
Economic
Dynamics
& Control

www.elsevier.nl/locate/econbase

Investment behavior under Knightian uncertainty – An evolutionary approach

Terje Lensberg*

Norwegian School of Economics and Business Administration, Helleveien 30, 5035 Bergen-Sandviken, Norway

Abstract

The 'as if' view of economic rationality defends the profit maximization hypothesis by pointing out that only those firms who act as if they maximize profits can survive in the long run. Recently, the problem of arriving at a logically consistent definition of rational behavior in games has shown that one must sometimes study explicitly the evolutionary processes that form the basis of this view. The purpose of this paper is to investigate the usefulness of genetic programming as a tool for generating hypotheses about rational behavior in situations where explicit maximization is not well defined. We use an investment decision problem with Knightian uncertainty as a borderline test case, and show that when the artificial agents receive the same information about the unknown probability distributions, they develop behavior rules as if they were expected utility maximizers with Bayesian learning rules. © 1999 Elsevier Science B.V. All rights reserved.

JEL classification: B41; C63; D83

Keywords: Knightian uncertainty; Bayesian rationality; Genetic programming

1. Introduction

Let S be a set of decision situations, and A a set of possible actions that can be taken in those situations. A *behavior theory* is a function

$$F : S \to A \qquad (1)$$

* Tel.: + 47-5595-9206; fax: + 47-5595-3939; e-mail: terje.lensberg@nhh.no.

PII: S0165-1889(98)00085-2

that associates an action to each possible situation. The set of situations may be, e.g. a class of games, in which case each element of A is a strategy, or S may be a class of investment problems under uncertainty, and the elements of A amounts invested.

In economics, the standard approach to constructing behavior theories is to impose axioms on the function F that describe intuitive notions of rationality, and then characterize the set of behavior theories that satisfy the axioms. Familiar examples are consumer theory, Bayesian decision theory and game theory, and a common theme in all of them is some form of explicit or implicit maximization hypothesis.

Recently, it has become clear that there are logical problems associated with extending the notion of Bayesian rationality to games: For example, Aumann and Brandenburger (1995) have shown that when a game has more than one Nash equilibrium, there is no guarantee that rational agents will end up in a Nash equilibrium at all. And if one considers games with an explicit time structure, there are a number of impossibility results that deny the existence of *any* logically consistent definition of rational play in such games, even for very weak intuitive notions of individual rationality (Basu, 1990; Bicchieri, 1989; Reny, 1992).

An alternative to the axiomatic approach was suggested by Alchian (1950), Friedman (1953) and Koopmans (1957), based on the idea of economic survival of the fittest, with profit maximization as an outcome of competition rather than a premise for it. Koopmans argued that

> If [survival] is the basis for our belief in profit maximization, then we should postulate that basis itself and not the profit maximization which it implies in certain circumstances … Such a change in the basis of economic analysis would … prevent us, for purposes of explanatory theory, from getting bogged down in those refinements of profit maximization theory which endow the decision makers with analytical and computational abilities and assume them to have information-gathering opportunities such as are unlikely to exist or be applied in current practice (Koopmans, 1957), (pp. 140–141).

A beautiful demonstration of the power of this approach has been given by Latané (1959), Breiman (1961) and Hakansson (1971): In a context of investment and capital accumulation under systematic risk, they show that in the long run, all capital will be held by those investors who act as if they maximize expected logarithmic utility, period by period. This result is interesting because it makes no assumptions about the preferences or motivation of individual agents, and yet produces extremely sharp predictions about aggregate behavior. Recently, Blume and Easley (1992) have generalized this approach and extended it to include the learning aspect of Bayesian rationality as well: Suppose the investors have subjective beliefs about the probability distributions that determine returns on investment, and suppose they maximize expected logarithmic utility given those beliefs. Then in the long run, all capital will be held by those investors who act as if they use a Bayesian learning rule to update their beliefs.

The recent influx of ideas from biology to economics has had a parallel in computer science, where it has produced a number of techniques for modeling artificially intelligent agents of bounded rationality. This toolbox is of great potential value to economists, who are concerned with closely related modeling problems. Two examples of successful applications, based on the pioneering work of Holland (1975), are Axelrod's (1987) use of a genetic algorithm to analyze evolution of strategies in the finitely repeated prisoners' dilemma, and Marimon et al. (1987) study of evolution of a general medium of exchange in a population of agents modeled as classifier systems.

A recent addition to this toolbox is *genetic programming* (GP) (Koza, 1992), which can be thought of as a technique for programming computers by natural selection. For our purpose, each program is a behavior rule F_i as defined in Eq. (1). The GP algorithm uses a large population of competing rules F_i whose behavior $F_i(s)$ is repeatedly computed for randomly selected situations $s \in S$ and evaluated to obtain a ranking of the rules in terms of fitness. Low-performing rules are replaced by genetic recombinations of high-performing ones, and the process continues until the whole population converges on some common behavior rule F, which is then proclaimed the outcome of the evolutionary process.

This author believes that GP will prove to be a useful tool for studying evolution of play in games with more than one Nash equilibrium. In this setting, there is a need for alternative hypotheses about rational behavior that can be subjected to theoretical investigation, and GP is a modeling technique which is capable of generating such hypotheses. However, because the issue of rationality is still unsettled for games with multiple Nash equilibria, it will be of interest to first investigate whether GP is able to generate rational behavior in situations where we know what rationality means. In this paper, we make an attempt in that direction by considering a borderline case, which shares with games the feature of leaving the agents in confusion about the probabilities they are facing, while still permitting us to recognize rational behavior when we see it. It is a version of Latané's (1959) investment decision problem, extended, as in Blume and Easley (1992), to allow for uncertainty about the relevant probability distributions. In contrast to Blume and Easley, we consider a situation with persistent Knightian (Knight, 1921) uncertainty, meaning that the agents have no a priori beliefs about returns on investments, and no learning rule that could be used to arrive at such beliefs. Our main result is that, despite all this bounded rationality, surviving agents act as if they knew the true prior, as if they used Bayes' rule to update it with respect to the available information on the current investment alternatives, and as if they maximized expected logarithmic utility given the posterior.

The remainder of the paper is organized as follows: In Sections 2 and 3, we describe the investment model, and give an outline of genetic programming in that context. Section 4 describes a typical run from the experiments, and Sections 5 and 6 presents the results of 20 independent runs, along with analyses of the deviations from Bayesian rational behavior. Section 7 concludes.

2. The investment model

Consider an economy operating during time periods $t = 1, 2, \ldots, \infty$, with a set $I := \{1, \ldots, m\}$ of *agents*, and two commodities; labor and a perishable consumer good. Depending on the course of events so far, an agent may own a firm, in which case he is a *capitalist*, or he may be an *entrepreneur* about to start one, or he may neither, in which case he is a *worker*. All agents supply one unit of labor in each period, and they prefer more goods to less.

Let C_t denote the set of capitalist-firms when period t begins. Each firm then owns a prepaid labor contract with one or more agents, and we denote by w_t^i the total amount of labor at the disposal of firm i. Each firm produces goods by employing a fraction $x_t^i \in [0, 1]$ of its labor force in a risky technology, which yields either 0 or 2 units of goods per unit labor input, and using the remaining part $1 - x_t^i$ in a riskless technology, which always yields 1 unit of goods per unit labor input. If we let $\tilde{\sigma}_t$ be a random variable which is $+1$ in the *good state* and -1 in the *bad state*, we may express the uncertain output $\tilde{q}_t^i$ of firm i in period t as

$$\tilde{q}_t^i = w_t^i(1 + \tilde{\sigma}_t x_t^i). \tag{2}$$

If the output of some firm is zero, it goes *bankrupt* and its owner becomes a worker. A worker may then choose to spend the next period as an entrepreneur, in which case he makes a commitment now to start a new firm next period by working full time in it during that period. The non-bankrupt firms trade their supplies of goods against labor contracts with the non-entrepreneurs for the next period, which then begins.

Let $q_t := \sum_{i \in C_t} q_t^i$ denote the aggregate supply of goods. The aggregate supply of labor is $m - e_t$, where e_t is the number of agents who decide in period t to be entrepreneurs next period. Assuming perfect competition, the price of goods in terms of labor is

$$\pi_t := (m - e_t)/q_t \tag{3}$$

and the amount of labor at the disposal of firm i in period $t + 1$ is

$$w_{t+1}^i := \pi_t q_t^i. \tag{4}$$

Letting B_t denote the set of capitalist-firms that go bankrupt in period t, and by E_t the set of entrepreneurs, the set of capitalist-firms in period $t + 1$ is given by

$$C_{t+1} = C_t \backslash B_t \cup E_t.$$

Moreover, using Eqs. (2)–(4), the amount of labor at the disposal of firm $i \in C_{t+1}$ in period $t+1$ can be written as

$$w_{t+1}^i = \begin{cases} (m - e_t)q_t^i/q_t & \text{if } i \in C_t \backslash B_t, \\ 1 & \text{if } i \in E_t. \end{cases} \tag{5}$$

This implies that $\sum_{i \in C_t} w_t^i = m$ for all t, hence w_t^i/m is firm i's share of the total wealth in the economy at time t.

Let $p_t := \Pr\{\tilde{\sigma}_t = 1\}$ denote the probability that the good state obtains in period t. We assume that the *success probabilities* p_t are generated by independent draws from a uniform probability distribution on $[0, 1]$. The firms, however, know neither the success probabilities nor the probability distribution from which they are drawn. All they observe is a random number of draws from the probability distribution p_t, of which g_t denotes the number of good outcomes and b_t denotes the number of bad ones.

The pair (g_t, b_t) is the *information* available to firm i when it makes its investment decision at time t. Note that all firms face the same (unobservable) success probabilities p_t and the same information (g_t, b_t), hence the environment is characterized by *systematic risk* and *symmetric information*.

If one were to solve the investment decision problem using Bayesian decision theory, one would equip each firm with a subjective prior probability distribution for the success probability p; a likelihood function for updating the prior with respect to new information, and a utility function u_i which is then maximized, given the available information. In general, the utility functions will depend on infinite sequences of payoffs, but in many cases of theoretical interest, one makes simplifying assumptions about risk and time preferences which allow one to reduce the intertemporal maximization problem to a sequence of myopic ones, and obtain a straightforward characterization of the optimal investment behavior rule: If we let $\hat{p}_t^i := \mathrm{E}_t^i[\tilde{p}_t \,|\, (g_t, b_t)]$ denote firm i's point estimate of the current success probability given its current prior and the available information (g_t, b_t), then its optimal investment ratio x_t^i is obtained as a function of w_t^i, g_t and b_t by solving

$$\max_{x_t^i} \hat{p}_t^i u_i(w_t^i(1 + x_t^i)) + (1 - \hat{p}_t^i)\, u_i(w_t^i(1 - x_t^i)). \tag{6}$$

Here, however, we will not assume that the firms solve maximization problems, only that they have *some* behavior rule $F_i : \mathbb{R}^3 \to [0, 1]$, where $x_t^i := F_i(w_t^i, g_t, b_t)$ is the investment ratio of firm i at time t, if it has information (g_t, b_t).

The mean behavior of the population will change over time as a result of (i) changes in the behavior rules F_i of individual firms, and (ii) changes in the wealth distribution. Changes in the wealth distribution are determined by Eq. (5), while changes in individual behavior rules arise through bankruptcy of existing firms, startups of new firms, and also *reorganizations* of existing firms. The development of individual behavior rules will be modeled as an evolutionary process by means of a genetic programming algorithm that we describe next.

3. The GP algorithm

There are many variants of GP algorithms, see Koza (1992) for an introduction, and Kinnear (1994) for a recent overview. Here we shall use a *steady-state* algorithm with *tournament selection*, which works along the following lines:

1. Set $t = 0$, and generate a population of m firms, each one equipped with an initial wealth of 1, and a randomly chosen behavior rule.
2. Set $t := t + 1$. Randomly select a success probability p_t from the interval $[0, 1]$, and use it to randomly generate the information (g_t, b_t). Calculate the investment ratio $x_t^i = F_i(w_t^i, g_t, b_t)$ of each firm, and update their wealth using Eq. (5).
3. Breed a number of new firms by replacing the behavior rules of unfit firms with genetic recombinations of the behavior rules of fit firms. Fitness is defined as accumulated wealth, and the genetic recombinations consist of a crossover operation involving two rules, and a mutation operation on one rule.
4. Go to 2 unless $t = t_{\max}$.

A behavior rule plays two different roles in the GP algorithm: On the one hand, it is a function which determines the behavior and fitness of the individual (the phenotype of the individual), and on the other, it has a tree structure which describes its properties from the point of view of genetic recombination (the genotype of the individual). Fig. 1 illustrates these two ways of looking at the same thing for the rule $(g - 1)/(w \cdot b)$. The figure also shows the rule expressed in LISP prefix notation, which captures the tree-structure of algebraic expressions better than the usual infix notation.

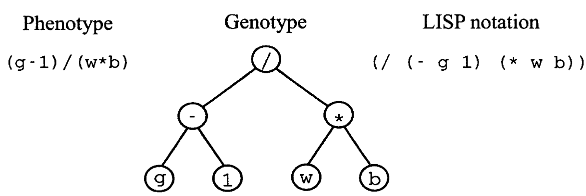

Fig. 1. Alternative representations of behavior rules.

In order to generate the behavior rules, evaluate them, and perform the genetic recombinations, the algorithm uses a number of basic elements that we describe next:

Terminal set: The terminal set consists of the observable variables and the constants of the problem. In our case, there are three observable variables, w^i, g and b, and a set $\Re$ of real constants.

Function set: These are the primitive functions of the problem, for which we take the four arithmetic operations $\{+, -, *, /\}$.

Fitness cases: The family of all investment decision problems S constitutes the set of fitness cases for the problem. Each such problem is a vector $(w^i, g, b; \mathrm{p})$, where w^i is a non-negative real number, g and b are non-negative integers, and where $\mathrm{p} \in [0, 1]$ is an unobservable success probability.

Fitness definition: In our context, the fitness of an individual firm is simply its accumulated wealth w^i_t.

Interpreter: This is a function $\mathscr{I}$ which operates on the response values of any behavior rule F_i in any situation $s \in S$ to yield an output $\mathscr{I}(F_i(s), s)$. This quantity is the experimenter's interpretation of the action taken by behavior rule F_i in situation s. In our context, we use the interpreter to constrain the actions to lie in the interval $[0, 1]$ by defining $\mathscr{I}$ as $\mathscr{I}(x, s) := \max[0, \min[1, x]]$.

Using these elements, one can define the following operations:

Rule initialization: The initial population of firms is generated by constructing m behavior rules at random. To construct the behavior rule of an individual firm, one starts with a randomly selected function and recursively builds a tree where the root of each subtree is a randomly selected function, and each leaf is a randomly selected terminal.

Mutation: To mutate a rule, one selects one of its subtrees at random and replaces it by a new subtree that is constructed from scratch in the manner just described.

Crossover: To cross two rules, one selects a random subtree in each of them and replaces the subtree of the first rule by the subtree of the second one. This operation guarantees that the offspring is a syntactically valid tree.

Breeding: To breed a new firm from the existing population of firms, one proceeds as follows:

- Randomly select three firms from the population and rank them in descending order according to their fitness.
- If firm 1 is not bankrupt, then replace the rule of firm 3 by a genetic recombination of the rules of firms 1 and 2. With high probability the recombination is a crossover of the rules of firms 1 and 2, and with low probability, it is a mutation of the rule of firm 1. If firm 1 is bankrupt, we replace the behavior rule of firm 3 by a new rule created from scratch.
- If firm 3 is bankrupt, the newly bred version of firm 3 is a *startup*, in which case it receives an initial wealth of 1. Otherwise, it is a *reorganization*, in which case its initial wealth is left unchanged.

Table 1
Main parameters of GP algorithm

Parameter	Value
PopulationSize	2000
NumberOfPeriods	100,000
FunctionSet	$(+, *, -, /)$
TerminalSet	$(w, g, b, \Re)$
RequiredTerminals	(g, b)
MaxNewRuleDepth	5
MaxRuleLength	100
BreedingRate	4
MutationProbability	0.1
CrossoverProbability	0.9

Table 1 lists the main parameters of the GP algorithm. We use a PopulationSize of 2000 with NumberOfPeriods set to 100,000. As mentioned earlier, each behavior rule is an arithmetic expression composed of functions from the FunctionSet $(+, *, -, /)$ of arithmetic operations, and of variables and constants from the TerminalSet $(w, g, b, \Re)$, where $\Re$ is the set $\{i/100.0 \mid i \in \{0, \dots, 100\}\}$ of real constants.

When a new behavior rule is generated from scratch, we repeatedly generate new candidates until we have found one in which the RequiredTerminals g and b are both present. This restriction is used in order to improve the genetic properties of the initial population somewhat, but is not imposed on behavior rules that are generated by crossover and mutation. For each candidate, we use uniform probability distributions to first select between the function set and the terminal set; second between the elements of each set, and third between each constant, if a random constant in $\Re$ was selected in step two.

The parameter MaxNewRuleDepth restricts the depth of new expression trees generated from scratch to be at most 5, with equal probabilities of generating rules of depths 2, 3, 4 and 5. MaxRuleLength restricts total number of functions and terminals in any expression tree to be no greater than 100. In each period, we use a BreedingRate of 4 to generate 4 new behavior rules by genetic recombination: With MutationProbability 0.1, we do a mutation, and with CrossoverProbability 0.9, we do a crossover.

To select nodes in an expression tree for crossover or mutation, we use a probability distribution where the probability of selecting a particular node is proportional to the number of functions and terminals in the subtree starting at the given node.

The GP algorithm basically works by maintaining a diverse population and subjecting it to selection pressure. The parameter values in Table 1 have therefore been chosen to enhance diversity while still permitting us to detect any emergent structure with some degree of precision. Larger values for PopulationSize, MaxNewRuleDepth, MaxRuleLength and BreedingRate will increase the diversity, and a substantial reduction in all these parameters might even cause a failure of the algorithm to detect rules that can avoid bankruptcy. On the other hand, a high BreedingRate will encourage growth of behavior rules which temporarily do well due to sheer luck. We have therefore chosen to use a low BreedingRate to keep this kind of noise at a minimum, and to compensate for the resulting slow convergence by using a large NumberOfPeriods.

The FunctionSet contains only primitive functions that turn out to be useful for solving the problem at hand. In preliminary experiments, we did include a number of other functions as well, which led to slower convergence without changing the qualitative results. Inclusion of irrelevant variables in the TerminalSet is known to have a similar effect (Koza, 1992). Our choice of a small MutationProbability and a large CrossoverProbability follows standard practice, although recent research (Angeline, 1997) indicates that mutation may in fact perform better than crossover for evolving expression trees of the type we use here.

We conclude this section with a description of the mechanism used to generate the fitness cases used by the GP algorithm. Recall that the fitness case faced by firm i at time t is a vector $(w_t^i, g_t, b_t; p_t)$, where p_t is an unobservable success probability and (g_t, b_t) is the observable information about p_t, and where w_t^i is the accumulated wealth of firm i when period t begins.

In order to generate the fitness case at time t, we begin by drawing a success probability p_t from the uniform distribution on $[0, 1]$. Next, we draw a realization n_t of the random variable $\tilde{n} := \mathrm{Int}(205/\tilde{u}) - 5$, where $\tilde{u}$ is another random variable which is uniformly distributed on the interval $[1, 41]$, and $\mathrm{Int}(z)$ is the nearest integer to z. Finally, we draw n_t realizations from the probability distribution p_t on the set $\{1, -1\}$ of good and bad outcomes, and denote by g_t and b_t the number of good and bad outcomes, respectively.

The probability distribution for $\tilde{n}$ has been chosen to focus on decision situations where the agents know very little about the unknown success probabilities. For example, there is some 50% probability that $\tilde{n} \leq 4$, and some 10% probability that $\tilde{n} = 0$, in which case the firms receive no information about the success probability p_t.

4. A successful behavior rule

In this section, we discuss the results of a first experiment with the GP algorithm on the investment problem described in Section 2. A typical example of a long-lived and highly fit behavior rule from the experiment is depicted in Fig. 2.

```
(+ (/ (- G B) (+ 0.72 (+ (+ B G) (+ 0.93 0.3))))
   (* (+ 0.72
         (/ (- G B)
            (+ (/ (- G B) (/ (- G B) (+ G 0.96)))
               (* (+ B
                     (+ (+ (* (* (* (+ (/ (- G B) (/ G 0.43)) 0.81) G) W)
                                 (- W W)) (+ (+ B G) W))
                        (+ (+ B G) (* (- G B) (+ G (- W (* 0.58 0.02)))))))
                  (- W W)))))
      (- W W)))
```

Fig. 2. A long-lived and highly fit behavior rule.

The behavior of this rule is completely determined by the first line of the expression. The rest is identically zero, and hence junk from a behavioral point of view.[1] The whole expression therefore simplifies to

$$\frac{g-b}{g+b+1.95}.$$

Using the interpreter $\mathscr{I}(\cdot)$ to restrict the response value of the behavior rule to the relevant interval $[0,1]$, we obtain the rule $F(w,g,b)$ defined by

$$F(w,g,b) := \max\left\{0, \frac{g-b}{g+b+1.95}\right\}. \tag{7}$$

There are several interesting features to note about the behavior of this rule: First, it never invests anything in the risky technology unless $g > b$, i.e. unless there is reason to believe that the success probability p is greater than $\frac{1}{2}$, and hence that the risky alternative yields a higher expected return than the riskless one. Second, a firm equipped with this behavior rule has zero probability of going bankrupt, since it never invests all its wealth in the risky alternative. And third, it invests the same fraction of its wealth in the risky alternative independently of its wealth.

These features are all representative of risk averse expected utility maximizers, and the third one strongly suggests constant relative risk aversion. It is therefore natural to investigate whether the GP algorithm has rediscovered the theoretical results of Latané (1959), Breiman (1961) and Hakansson (1971), that long-run survival in a situation with systematic risk implies maximization of expected logarithmic utility, period by period.

In our model, one-period expected logarithmic utility is given by

$$p\,\log(w(1+x)) + (1-p)\log(w(1-x)),$$

[1] The long zero term in lines 2–10 is quite useful from the *reproductive* point of view, however: It increases the probability that a mutation or crossover will leave line 1 intact and yield an offspring with the same behavior as the highly fit parent.

cf. Eq. (6), and maximization with respect to x yields the behavior rule

$$F^*(\mathrm{p}) := \max\{0, (2p - 1)\}. \tag{8}$$

However, since our artificial agents operate under Knightian uncertainty, they do not know p, only the imperfect signal (g, b). The results of Latané, Breiman and Hakansson do not cover this case, but if the agents were clever enough, they might eventually be able to figure out that the success probabilities p_t are drawn from a uniform distribution and that p_t therefore has a Beta distribution with parameters $n' = g + b + 2$ and $r' = g + 1$. A Bayesian rational agent would then conclude that

$$\mathrm{E}[\tilde{p}] = \frac{g + 1}{g + b + 2}, \tag{9}$$

and substitute $\mathrm{E}[\tilde{p}]$ for p in the Bayesian rational behavior rule (8). This yields

$$F^*(g, b) = \max\left\{0, \frac{g - b}{g + b + 2}\right\}, \tag{10}$$

which is almost identical to the behavior rule (7) that was generated by the GP algorithm.

A firm equipped with this behavior rule will behave in a manner which is not only consistent with expected utility maximization, it will also act as if it were familiar with Bayesian statistics, despite the fact that it has no way of understanding the nature of the investment problem and no apparatus for estimating the relevant parameters of the unknown probability distributions. We thus have an example which shows that GP is capable of generating behavior rules that are almost Bayesian rational.

5. Results from 20 GP-runs

In order to investigate the robustness of this result, we did 20 runs with the GP algorithm, and collected data that describe the behavior of each population. Since there is a steady inflow of new firms into the population with a variety of behavior rules that for the most part survive for only a small number of periods, we have included in our statistical measures only those who have survived for 500 periods or more. For every period t and run r, the mean behavior of the population at information state (g, b) is defined as

$$F_{rt}(g, b) := \sum_{i \in V_{rt}} \alpha^i_{rt} F^i_{rt}(w^i_{rt}, g, b),$$

where for each time t and run r, V_{rt} is the set of firms that have survived for at least 500 periods, F^i_{rt} is the behavior rule of firm i, w^i_{rt} is its wealth, and $\alpha^i_{rt} := w^i_{rt} / \sum_{j \in V_{rt}} w^j_{rt}$.

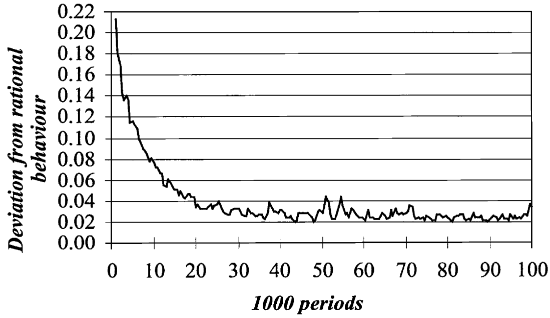

Fig. 3. Standard deviation from Bayesian rational behavior across all rounds over intervals of 500 periods.

We next define the total mean square deviation from Bayesian rational behavior in period t of run r as

$$\mathrm{TMS}_{rt} := \sum_{i \in V_{rt}} \alpha^i_{rt}[F^i_{rt}(w^i_{rt}, g_{rt}, b_{rt}) - F^*(g_{rt}, b_{rt})]^2,$$

where F^* is the Bayesian rational behavior rule defined in Eq. (10), and (g_{rt}, b_{rt}) is the information faced by firm i about the unknown success probability p_{rt}.

Fig. 3 depicts the development of the standard deviation from Bayesian rational behavior, calculated for every 500 periods as the square root of the mean of TMS_{rt}, taken across the recent 500 periods and all 20 rounds.

The figure shows that the deviation from Bayesian rational behavior is reduced to its minimum during the first 30.000 periods. From then on, it fluctuates within the range 0.02–0.04, with an occasional peak. These peaks occur when some behavior rule with low-risk aversion hits a streak of good luck and becomes the all-dominating firm of its population for a limited number of periods.

It turns out that on average across all rounds, only some 30% of the total deviation from Bayesian rational behavior is due to deviation by the mean behavior of the population. The remaining 70% is unsystematic variation due to deviation by individual firms from the mean behavior of the population. This is illustrated in Fig. 4, which depicts the total deviation from Bayesian rational behavior and its variance decomposition over the last 25.000 periods for each of the 20 runs.

As can be seen from the figure, the aggregate behavior across the last 25.000 information states deviates from the Bayesian rational behavior by less than 0.01 in most runs. Since for each information state, the rational behavior is a real number in the interval $[0, 1)$, we conclude that the aggregate behavior of the population is almost Bayesian rational for most states.

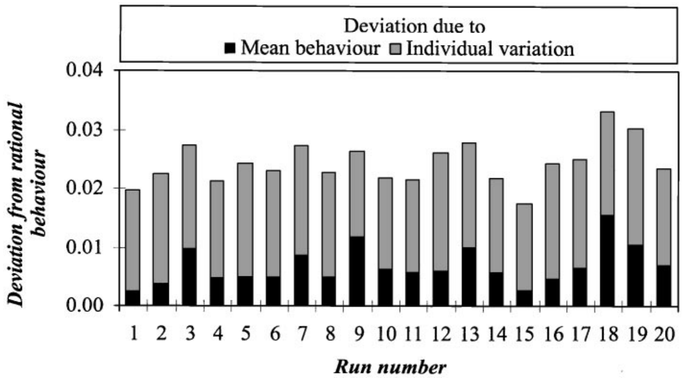

Fig. 4. Variance decomposition of the total deviation from Bayesian rational behavior for individual runs over last 25.000 periods.

6. Uncertainty attitude

We next investigate whether there is any systematic variation in the mean deviation from Bayesian rational behavior across different information states. Of particular interest is whether our artificial agents are *uncertainty averse* in the sense of Ellsberg (1961). Roughly speaking, an agent is uncertainty averse if she prefers known probability distributions to unknown ones, and there is a large body of experimental evidence that human decision makers tend to display this phenomenon (see Camerer and Weber, 1992 for a recent overview of the literature).

However, since expected utility is linear in the probabilities, uncertainty aversion is inconsistent with Bayesian rationality. In particular, a Bayesian rational investor does not care whether the success probability p_t is objectively known or her subjective expectation of a random variable $\tilde{p}_t$ (cf. Eqs. (9) and (10)). In contrast, an uncertainty averse investor would prefer to invest less in the risky alternative if she has less information about $\tilde{p}_t$ (smaller n_t), even if $\mathrm{E}[\tilde{p}_t \mid g_t, b_t]$ is constant (Dow and Werlang, 1992).

In order to test the uncertainty attitude of the developed behavior rules, we calculated for every 500th period t in every run r, the mean behavior of the population at information state (g, b), for a set of information states such that $\mathrm{E}[\tilde{p} \mid (g, b)] = \frac{2}{3}$. As can be seen from Eq. (9), this is the case when $g = 2b + 1$, which yields a Bayesian rational investment ratio of $\frac{1}{3}$. Fig. 5 shows the mean behavior of the population for a set of such information states across all runs over the last 25.000 periods, together with a 90% confidence band. The confidence band is calculated for each (g, b) using the mean behavior of each run over the last 25.000 periods as one observation, which yields a total of 20 observations.

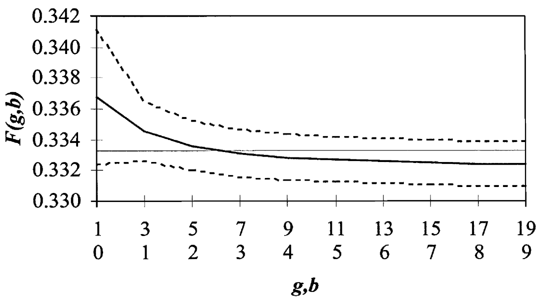

Fig. 5. Mean behavior over last 25.000 periods with 90% confidence band for information states (g, b) such that $F^*(g, b) = \frac{1}{3}$.

The figure shows that the firms tend to be a little too bold when they only have a small amount of information, and a bit too careful when they have more, as compared to the Bayesian rational investment ratio, which is $\frac{1}{3}$ for each information state in Fig. 5. This is exactly the opposite of the uncertainty averse behavior found by Ellsberg (1961). Thus while real people tend to prefer known probability distributions to unknown ones, our artificial agents seem to have the opposite preference. As can be seen from the figure, this deviation from the Bayesian rational behavior is not statistically significant in the data set considered here. However, the phenomenon seems to be quite robust across many different data sets and with different variants of the GP algorithm, and it is therefore natural to look for possible explanations.

Some insight can be gained by studying how the firms' uncertainty attitude changes over time. To this end, we calculated their uncertainty attitude at information state $(g, b) = (1, 0)$, by comparing $F(1, 0)$ to $F(3, 1)$, where F is the mean behavior of the population. If $F(1, 0) - F(3, 1)$ is positive, zero or negative, then F is said to be *uncertainty prone*, *uncertainty neutral*, and *uncertainty averse*, respectively.

Fig. 6 depicts the development of $F(1, 0) - F(3, 1)$, calculated across all 20 runs, and across 20 sets of periods of 5.000 periods each. The dotted lines represent a 90% confidence band calculated for each set of periods using the mean behavior of each run across sets of 5.000 periods as one observation, which yields a total of 20 observations.

The figure shows that the firms are extremely uncertainty prone during the first 10.000 periods or so of each run. However, there is considerable variation across runs until uncertainty neutral firms begin to take over during the next 10.000 periods. During the last 75.000 periods, the initial uncertainty proneness has largely disappeared. Most likely, however, the initial uncertainty proneness is still represented in the genes of the behavior rules in later periods, even if it

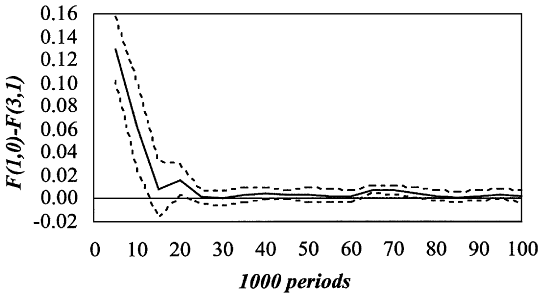

Fig. 6. Mean uncertainty attitude $F(1,0) - F(3,1)$ across all runs for sets of 5.000 periods.

does not necessarily affect their actual behavior.[2] As long as this genetic material is present in the population, it will continue to produce some behavior rules which may be responsible for the weak tendency to uncertainty proneness in later periods.

Of course, this raises the question of why the firms are so uncertainty prone in the early periods. To see why, we listed for each run the most fit behavior rule in periods $\{5500, 6000, \dots, 15{,}000\}$, and looked for common characteristics of these 200 rules. As it turned out, more than half of them had an arithmetic structure identical to one of the following:

$$F_1(g,b) = \frac{g-b}{g+g}, \qquad F_2(g,b) = \frac{g-b}{g/k}, \qquad F_3(g,b) = \frac{g-b}{g+k},$$

where k is one of the constants in $\Re$, i.e. a real number between 0 and 1. In our sample of 200 most fit rules, the constant k was typically in the neighborhood of 0.5 for rules of type F_2, and close to 1 for rules of type F_3. Note that all rules of these types are uncertainty prone, in the sense defined earlier.

One reason why these three types of rules are so dominant initially, may be that their ratio of efficiency to complexity is very high: The numerator $g - b$ is a precise signal of *when* it is profitable to invest in the risky alternative, and the role of the denominator is to determine *how much* to invest. Since the whole expression is so small, there is a fairly small number of such expressions, and therefore they have a fairly high probability of being generated from scratch or by genetic recombination.

Nevertheless, there are a number of other behavior rules with this structure, and the question remains why none of them are represented in our sample of most fit rules.

[2] The behavior rule with the large zero term in Fig. 2 illustrates how this might occur.

With our sets of functions and terminals, there is a total of 34 distinct behavior rules with a numerator of $g - b$ and a denominator consisting of 3 or fewer functions and terminals. Among those 34 rules, there are only four that never face a risk of going bankrupt. One of these four rules always yields an output of zero, and hence is completely useless, since it never invests anything in the risky alternative, no matter how profitable it is to do so. The remaining three rules are exactly the ones that were represented by more than 50% in our sample of most fit behavior rules.

We therefore conclude that the reason why these three rules tend to dominate the populations early on, is a combination of simplicity, efficiency and zero bankruptcy risk, and that the uncertainty proneness which we observe in the early periods is just a by-product of the evolutionary pressure which selects in favor of behavior rules with these three features. However, the uncertainty proneness does impose a cost in terms of lost profits on the firms that host these rules, and in later periods they are replaced by firms with more rational behavior rules.

7. Concluding remarks

The purpose of this paper has been to investigate the usefulness of genetic programming as a tool for generating hypotheses about rational behavior in situations where the issue of rationality is not clear-cut. To this end, we have tested it on the borderline case of Knightian uncertainty, where the agents do not know the true probability distributions or how they are generated, but where it is still possible to recognize rational behavior when one sees it.

We have shown that the algorithm systematically generates behavior which is Bayesian rational, despite the fact that the artificial agents do not solve maximization problems and have no apparatus for estimating the relevant parameters of the unknown probability distributions. Although GP does not produce theorems, the approach has a number of interesting features that we would like to summarize at this point.

Observe first that GP yields results at the same level of generality as a conventional analysis based on profit or utility maximization. In both cases, the outcome of the analysis is a function F which associates an action $F(s)$ to each possible decision situation $s \in S$ in which the decision maker might find himself. In a conventional analysis, the function F would be derived from the first-order conditions of one or more maximization problems, while in GP, it is a result of evolution, but the structure of the results from the two approaches is identical.

A key aspect of the approach is that individual agents are modeled as rigid rule followers who do not change their behavior over time, except through bankruptcy or reorganization. As in the evolutionary game theory of Maynard

Smith (1982), all change in behavior takes place at the level of the population, but GP differs from evolutionary game theory by not requiring the modeler to specify all possible behavior rules explicitly ex ante. In GP, new behavior rules emerge as a result of random recombinations of behavior rules that have been successful in the past, somewhat like Schumpeter's (1942) view of innovation.

The third point we would like to mention is that the behavior rules produced by GP are the result of learning by example, and generalization to an infinite space of decision situations from experience with a finite number of situations. Since the process of learning and knowledge generalization takes place at the level of the population, GP differs from other models of knowledge generalization, e.g. case based reasoning (Gilboa and Schmeidler, 1995), where the focus is on the learning process of the individual agent.

So what are the features of GP that allow it to represent knowledge generalization? Generally speaking, GP is a procedure for conducting a global search for the best alternative in a function space, where a finite number of points in the domain are used to test the quality of the candidates, and where generalization to the whole domain is achieved if the winning candidate is a continuous function. The potential area of economic applications for GP algorithms is therefore much wider than the one considered here. For example, a key element in the problem of computing equilibria in models with many heterogeneous agents (Ríos-Rull, 1997) is also to find a Bayesian rational behavior rule in a space of functions on a domain which represents possible variations in economic state and individual characteristics, and it might therefore be of interest to try out the technique for that type of problem as well.

In this paper, our interest in Knightian uncertainty has been motivated by the fact that this setting is similar to game-playing situations in that the agents do not know what probability distributions they are up against, especially in games with more than one Nash equilibrium. This is an area where the issue of rationality is still open, and where there is a need for alternatives to the existing models. It is therefore a natural topic for further research to study rational behavior in games, using the same technique that was shown to generate rational behavior under Knightian uncertainty.

Acknowledgements

I am grateful to Steinar Ekern, Heine Rasmussen and Hans Werner Sinn for helpful discussions and comments to earlier drafts. Thanks are also due a referee of this journal and to the referees and participants of the GP97 and CEF97 conferences for valuable criticism and suggestions. Remaining errors are mine. The research reported here was done in part during a visit at the Center for Economic Studies, University of Munich, and financial support from CES is gratefully acknowledged.

References

Alchian, A.A., 1950. Uncertainty, evolution and economic theory. Journal of Political Economy 58, 211–221.

Angeline, P.J., 1997. Subtree crossover: Building block engine or macromutation? In: e.a. Koza, J.R. (Ed.), Genetic Programming 1997. Proceedings of the Second Annual Conference. Morgan Kaufmann Publishers, San Fransisco, pp. 9–17.

Aumann, R., Brandenburger, A., 1995. Epistemic conditions for Nash equilibrium. Econometrica 63, 1161–1180.

Axelrod, R., 1987. The evolution of strategies in the iterated prisoner's dilemma. In Davis, L. (Ed.), Genetic Algorithms and Simulated Annealing. Pittman, London.

Basu, K., 1990. On the non-existence of a rationality definition for extensive form games. International Journal of Game Theory 19, 33–44.

Bicchieri, C., 1989. Self-refuting theories of strategic intercation: a paradox of common knowledge. Erkenntnis 30, 69–85.

Blume, L., Easley, D., 1992. Evolution and market behavior. Journal of Economic Theory 58, 9–40.

Breiman, L., 1961. Optimal gambling systems for favorable games. In: Neyman, J. (Ed.), Proceedings of the Fourth Berkeley Symposium on Mathematical Statistics and Probability. University of California Press, Berkeley, pp. 65–78.

Camerer, C., Weber, M., 1992. Recent development in modeling preferences: Uncertainty and ambiguity. Journal of Risk and Uncertainty 5, 325–370.

Dow, J., Werlang, S.R.C., 1992. Uncertainty aversion, risk aversion, and the optimal choice of portfolio. Econometrica 60 (1), 197–204.

Ellsberg, D., 1961. Risk, ambiguity, and the Savage axioms. Quarterly Journal of Economics 75, 643–669.

Friedman, M., 1953. Essays in Positive Economics. University of Chicago Press, Chicago.

Gilboa, I., Schmeidler, D., 1995. Case-based decision theory. Quarterly Journal of Economics 110, 605–639.

Hakansson, N.H., 1971. Capital growth and the mean–variance approach to portfolio selection. Journal of Financial and Quantitative Analysis 6, 517–557.

Holland, J.H., 1975. Adaptation in Natural and Artificial Systems. University of Michigan Press, Ann Arbor.

Kinnear, K.E., 1994. Advances in Genetic Programming. MIT Press, Cambridge.

Knight, F.H., 1921. Risk, Uncertainty and Profit. Houghton Mifflin, Boston.

Koopmans, T.C., 1957. Three Essays on the State of Economic Science. McGraw-Hill, New York.

Koza, J.R., 1992. Genetic Programming: On the Programming of Computers by Means of Natural Selection. MIT Press, Cambridge.

Latané, H.A., 1959. Criteria for choice among risky assets. Journal of Political Economy 67, 144–155.

Marimon, R.E., McGrattan, E., Sargent, T.J., 1987. Money as a medium of exchange in an economy with artificially intelligent agents. Journal of Economic Dynamics and Control 14, 329–373.

Maynard Smith, J., 1982. Evolution and the Theory of Games. Cambridge University Press, Cambridge.

Reny, P.J., 1992. Rationality in extensive-form games. Journal of Economic Perspectives 6, 103–118.

Ríos-Rull, J.-V., 1997. Computation of equilibria in heterogeneous agent models. Staff Report 231, Federal Reserve Bank of Minneapolis.

Schumpeter, J.A., 1942. Capitalism, Socialism and Democracy. Harper, New York.