CES Working Paper Series

INVESTMENT BEHAVIOUR UNDER KNIGHTIAN UNCERTAINTY - AN EVOLUTIONARY APPROACH

Terje Lensberg

Working Paper No. 126

1997

Center for Economic Studies
University of Munich
Schackstr. 4
80539 Munich
Germany
Telephone & Telefax:
++49-89-2180-3112

The research reported here was done in part during a visit at the *Center for Economic Studies*, University of Munich. Financial support from *CES* is gratefully acknowledged, and thanks are due to Hans Werner Sinn for his hospitality during my stay there.

CES Working Paper No. 126
January 1997

INVESTMENT BEHAVIOUR UNDER KNIGHTIAN UNCERTAINTY - AN EVOLUTIONARY APPROACH

Abstract

The paper analyzes investment behaviour under Knightian uncertainty by means of a genetic programming algorithm. This is an experimental approach which yields analytical results at a level of generality comparable to that obtained by conventional methods. When the artificial agents receive the same information about the unknown probability distributions, they develop behaviour rules as if they were expected utility maximizers with Bayesian learning rules and logarithmic utility functions. We then introduce asymmetric information, and study how it affects the agents' implicit preferences for risk and uncertainty.

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

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 *behaviour theory* is a function

$$F : S \to A \tag{1}$$

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, in which case the elements of A are amounts invested.

In economics, the standard approach to constructing behaviour theories is to impose axioms on the function F that describe intuitive notions of rationality, and then characterize the set of behaviour 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.

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 beatiful 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 behaviour. These results have recently been generalized to situations involving both systematic and idiosyncratic risk by Robson (1996) in an evolutionary framework.

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 modelling artificially intelligent agents of bounded rationality. This toolbox is of great potential value to economists, who are concerned with closely related modelling 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 of evolution of strategies in the finitely repeated prisoners' dilemma, and Marimon, McGrattan and Sargent's (1987) study of evolution of a general medium of exchange in a population of agents modelled 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 behaviour rule F_i as defined in (1). The genetic programming algorithm uses a large population of competing rules F_i whose behaviour $F_i(s)$ 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 behaviour rule F, which is then proclaimed the outcome of the evolutionary process.

The aim if the present paper is to investigate the usefulness of GP as an aid to studying rational behaviour in settings where the issue of rationality is not clearcut, and we approach this aim by considering a borderline case: It is a version of the investment decision problem of Latané (1959), Breiman (1961) and Hakansson (1971), extended to allow for Knightian (1921) uncertainty, meaning that the agents do not know the relevant probability distributions.

In Bayesian decision theory, the problem of Knightian uncertainty is dealt with by supplying the agents with a subjective a priori probability distribution, a likelihood function to be used for updating the a priori distribution with respect to new information, and a von Neumann–Morgenstern utility function which is then maximized given the available information. The result is a behavour theory which will depend on the parameters of subjective probability distributions, likelihood functions and utility functions. We show here that GP produces results that are consistent with Bayesian rationality, but more precise in the sense of leaving fewer parameters undetermined.

We believe that this type of results will be of interest for the following reason: If GP systematically generates rational behaviour in decision situations where we know what rationality means, then it is a potentially valuable tool for generating hypotheses about rational behaviour in situations where the issue of rationality is less clear–cut, as in games with more than one Nash–equilibrium.[1] The setting of investment decisions under Knightian uncertainty has been chosen because it shares with games the feature of leaving the agents in confusion about the probabilities they are facing, while still permitting them to be analysed by conventional methods.

[1] As shown by Aumann and Brandenburger (1995), it not clear that rational agents will end up in a Nash equilibrium if the game has more than one of them.

The remainder of the paper is organized as follows: In section 2, we describe the investment model, and in section 3, we give an outline of genetic programming in our context. Section 4 contains the results, and section 5 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; labour 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 labour 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 labour contract with one or more agents, and we denote by w_t^i the total amount of labour at the disposal of firm i. Each firm produces goods by employing a fraction $x_t^i \in [0, 1]$ of its labour force in a risky technology, which yields either 0 or 2 units of goods per unit labour input, and using the remaining part $1 - x_t^i$ in a riskless technology, which always yields 1 unit of goods per unit labour input. If we let $\tilde{\sigma}_t$ be a stochastic 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 labour 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 labour 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 labour is

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

and the amount of labour 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 entrpreneurs, it follows that the set of capitalist–firms in period $t + 1$ is

$$C_{t+1} = C_t \backslash B_t \cup E_t \tag{5}$$

and that the amount of labour at the disposal of firm $i \in C_{t+1}$ in period $t+1$ is given by

$$w_{t+1}^i = \begin{cases} (m - e_t)q_t^i/q_t & \text{if } i \in C_t \backslash B_t \\ 1 & \text{if } i \in E_t. \end{cases} \tag{6}$$

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 n_t 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.

A behaviour rule for firm i is then a function $F_i(w, g, b)$, where the *investment ratio* $x_t^i := F_i(w_t^i, g_t^i, b_t^i)$ is the fraction of firm i's wealth w_t^i which it employs in the risky technology at time t, if it has information (g_t^i, b_t^i).

The mean behaviour of the population will change over time as a result of (i) changes in the behaviour rules F_i of individual firms, and (ii) changes in the wealth distribution. Changes in the wealth distribution are determined by (6), while changes in individual behaviour rules arise through bankruptcy of existing firms, startups of new firms, and also *reorganizations* of existing firms. The development of individual behaviour rules will be modelled as an evolutionary process by means of a genetic programming (GP) algorithm that we describe next.

3 The genetic programming algorithm

In order to describe the GP algorithm, we begin with a broad overview and continue with more detailed descriptions of the basic elements and operations involved. There are many variants of GP algorithms, see Kinnear (1994) for a recent overview. Here we shall use a so-called *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 behaviour rule.

2. Set $t := t+1$. Randomly select a success probability p_t from the interval $[0, 1]$, and generate the information available to each firm. Calculate the investment made by each firm, and update their wealth using (6).

3. Breed a number of new firms by replacing the behaviour rules of unfit firms with genetic recombinations of the behaviour 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 behaviour rule plays two different roles in the GP algorithm: On the one hand, it is a function which determines the behaviour 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). Figure 1 illustrates these two ways of looking at the same thing for the rule $(g - 1)/(w * b)$. 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.

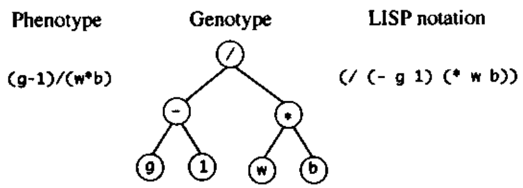

Figure 1: Alternative representations of behaviour rules

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

Terminal set The terminal set consists of the variables and constants of the problem. In our case, there are three variables, w, 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, g, b), where w is a non-negative real number, and where g and b are non-negative integers.

Fitness definition In our context, the fitness of an individual firm is simply its accumulated wealth w_t^i.

Interpreter This is a function $\mathcal{I}$ which operates on the response values of any behaviour rule F_i in any situation $s \in S$ to yield an output $\mathcal{I}(F_i(s), s)$. This quantity is the experimenter's interpretation of the action taken by behaviour 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 $\mathcal{I}$ as $\mathcal{I}(x, s) := \max[0, \min[1, x]]$.

Using these elements, one can define the following operations:

Rule initialization The initial population of firms is generated by constructing m behaviour rules at random. To construct the behaviour 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 3 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 behaviour 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.

The following table describes the main parameters used by the GP algorithm.

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

The GP algorithm is run with a PopulationSize of 2000 with NumberOfPeriods set to 100,000. As mentioned earlier, each behaviour rule is is an arithmetic expression composed of functions from the FunctionSet $(+, *, -, /)$ of arithmetic operations, and of variables and constants from the terminals from the **TerminalSet** $(w, g, b, \Re)$, where $\Re$ is the set $\{i/100.0 \mid i \in \{0, \ldots, 100\}\}$ of real constants.

When a new behaviour 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.[2]. 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 $\Re$ was selected in step two.

The parameter **MaxNewRuleDepth** restricts the depth of 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 behaviour 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.

We conclude this section with a description of the mechanism used to generate information about the success probabilities p_t. In each period t, we begin by drawing a success probability p_t from the uniform distribution on $[0,1]$. Next, we draw 200 realizations of good and bad outcomes from the probability distribution defined by p_t, and arrange them in a vector which we denote by H_t and think of as the *maximal* amount of information available at time t.

In order to generate some information for firm i, we first draw a realization n_t^i 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. The information given to firm i is then the pair (g_t^i, b_t^i), which consists of the number of good and bad outcomes among the first n_t^i elements of the vector H_t.

The probability distribution for the amount $\tilde{n}$ of information is shown in figure 2 for the first 20 values of n. As can be seen from the figure, we have chosen to focus on decision situations where the agents know very little about the unknown success probabilities. For example, there is some 50 per cent probability that $\tilde{n} \leq 4$, and some 10 per cent probability that $\tilde{n} = 0$, in which case the firms receive no information about the success probabability p_t.

[2]We use this restriction in order to improve the genetic properties of the initial population somwhat, but do not impose it on behaviour rules that are generated by crossover and mutation.

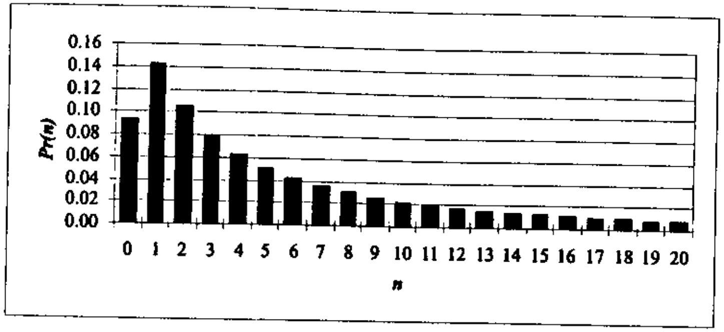

Figure 2: Probability distribution of $\tilde{n}$

4 Results

In this section, we report the results of two experiments with the GP algorithm on the investment problem described in section 2. Two cases will be considered: In *case 1*, all firms have access to the same information every period, and in *case 2* they receive different amounts of information.

4.1 Case 1: Equal information

A typical example of a long-lived and highly fit behaviour rule from our experiments with case 1 is the following:

```
(+ (/ (- 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)))
```

The behaviour of this rule is completely determined by the first line of the expression. The rest is identically zero, and hence junk from a behavioural point of view.[3] The whole

[3]The long zero term in lines 2-9 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 behaviour as the highly fit parent.

expression therefore simplifies to

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

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

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

There are several interesting features to note about the behaviour 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 $1/2$, and hence that the risky alternative yields a higher expected return than the riskless one. Second, a firm equipped with this behaviour 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 for 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

$$p\log(w(1+x)) + (1-p)\log(w(1-x)), \tag{8}$$

and maximization with respect to x yields the behaviour rule

$$F^*(p) := \max\{0, (2p-1)\}. \tag{9}$$

However, since our firms 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 our firms were indeed Bayesian rational, they would 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

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

and substitute $E[\tilde{p}]$ for p in the optimal behaviour rule (9). This yields

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

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

In order to investigate the robustness of this result, we did 20 runs with the GP algorithm, and collected data that describe the behaviour of each population. Since there is a steady inflow of new firms into the population with a variety of behaviour rules that for the most part only survive for 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 behaviour of the population at information state (g, b) is defined as

$$F_{rt}(g,b) := (1/w_{rt}) \sum_{i \in V_{rt}} w^i_{rt} F^i_{rt}(w^i_{rt}, g, b), \tag{11}$$

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 behaviour rule of firm i, w^i_{rt} is its wealth, and $w_{rt} = \sum_{i \in V_{rt}} w^i_{rt}$.

We next define the total mean square deviation from optimal behaviour in period t of run r as

$$\text{TMS}(r,t) := (1/w_{rt}) \sum_{i \in V_{rt}} w^i_{rt} \left[F^i_{rt}(w^i_{rt}, g^i_{rt}, b^i_{rt}) - F^*(g^i_{rt}, b^i_{rt})\right]^2 \tag{12}$$

where F^* is the optimal behaviour as defined defined in (10), (g^i_{rt}, b^i_{rt}) is the information faced by firm i about the unknown success probability p_{rt}.

Figure 3 depicts the development of the standard deviation from optimal behaviour, calculated for every 500th period as the square root of the mean of TMS(r,t), taken across the last 500 periods and across all 20 rounds.

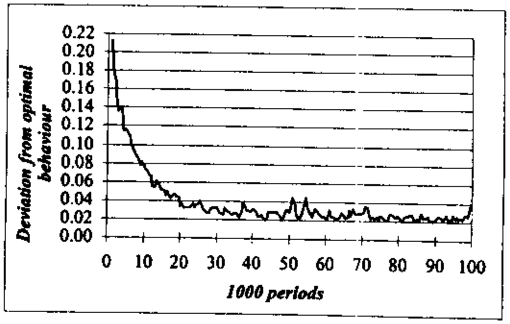

Figure 3: Standard deviation from optimal behaviour across all rounds over intervals of 500 periods (Equal information)

The figure shows that the deviation from optimal behaviour is reduced to its minimum during the first 30.000 periods. From then on, it fluctuates within the range 0.02 to 0.04, with an occasional peak. These peaks occur when some behaviour rule with low risk aversion hits a strike 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 per cent of the total deviation from optimal behaviour is due to deviation by the mean behaviour of the population from the optimal behaviour. The remaining 70 per cent is unsystematic variation due to deviation by individual firms from the mean behaviour of the population. This is illustrated in figure 4, which depicts the total deviation from optimal behaviour and its variance decomposition over the last 25.000 periods for each of the 20 runs.

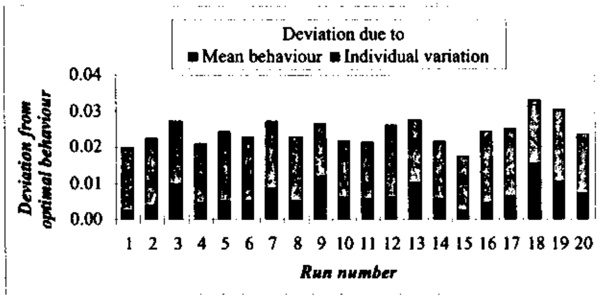

Figure 4: Variance decomposition of the total deviation from optimal behaviour for individual runs over last 25.000 periods (Equal information)

Apparently, the aggregate behaviour of the population is quite close to the optimal one for most runs. To see if there were any systematic differences, we calculated for every 500th period t in every round r, the mean behaviour of the population at information state (g, b), for a set of information states such that the optimal investment ratio is 1/3. As can be seen from (10), this is the case when $g = 2b + 1$. Figure 5 shows the mean behaviour of the population for a set of such information states across all runs over the last 25.000 periods, together with a 90 per confidence band. The confidence band is calculated for each (g, b) using the mean behaviour of each run over the last 25.000 periods as one observation, which yields a total of 20 observations.

As can be seen from the figure, the largest deviations from the optimal behaviour occur for small amounts of information. This is somewhat contrary to what one might expect apriori, since the firms have more experience with low-information states than with high-information states. 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 optimal investment ratio, which is 1/3 for each information state in figure 5.

This type of behaviour is exactly the opposite of the *uncertainty averse* behaviour found by Ellsberg (1961) and others in experiments with human decision makers: Real people tend to prefer known probability distribution to unknown ones, while our artificial agents seem to have the opposite preference. As can be seen from the figure, this

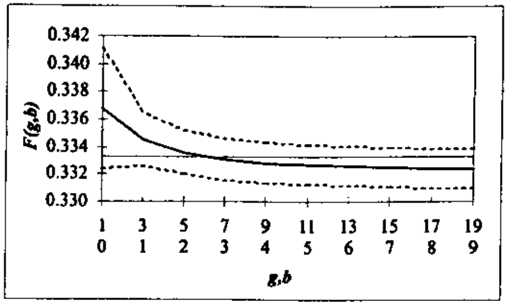

Figure 5: Mean behaviour over last 25.000 periods with 90 per cent confidence band for information states (g, b) such that $F^*(g, b) = 1/3$. (Equal information)

deviation from the Bayesian rational behaviour 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 differents variants of the GP algorithm, as long as the agents have equal information, 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 study their uncertainty attitude at information state $(g, b) = (1, 0)$, by comparing $F(1, 0)$ to $F(3, 1)$, where F is the mean behaviour 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.

Figure 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 per cent confidence band calculated for each set of periods using the mean behaviour of each run over the last 5.000 periods as one observation, which yields a total of 20 observations.

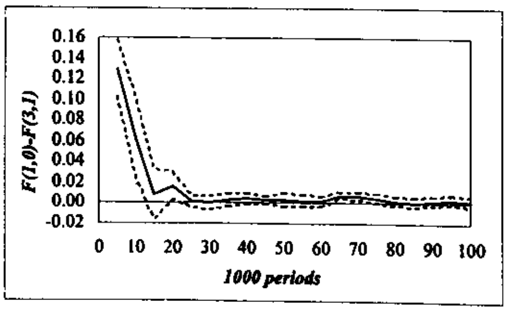

Figure 6: Mean uncertainty attitute $F(1, 0) - F(3, 1)$ across all runs for sets of 5.000 periods. (Equal information)

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 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 behaviour rules in later periods, even if it does not necessarily affect their actual behaviour.[4]. As long as this genetic material is present in the population, it will continue to produce some behaviour 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 behaviour rule in periods $\{5500, 6000, \ldots, 15000\}$, 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

$$F_1(g,b) = \frac{g-b}{g+g} \qquad F_2(g,b) = \frac{g-b}{g/k} \qquad F_3(g,b) = \frac{g-b}{g+k}, \tag{13}$$

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 neighbourhood 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.

The reason why these 3 types of rules are so dominant initially, is 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 behaviour rules with this structure, and the question remains why none of them are represented in our sample of most fit rules.

With our sets of functions and terminals, there is a total of 34 distinct behaviour rules with a numerator of $g - b$ and a denominator consisting of 3 or fewer functions and terminals. They are listed in table 1, along with sets of values for the parameters g and b for which the rules face a positive risk of bankruptcy.[5]

The table shows that among the 34 alternatives, there are only 4 rules that never face a risk of going bankrupt. Of these 4 rules, number 10 is identically equal to -1, and hence is completely useless, since it never invests anything in the risky alternative, no

[4]The behaviour rule with the large zero term on page 8 illustrates how this might occur.

[5]The routine for evaluating behaviour rule expressions uses an extended real number algorithm in order to avoid system crashes as a result of e.g. division by zero. For example, $x/0 = \text{sign}(x) \cdot \infty$ and $0/0 = -\infty$.

Rule	Behaviour	Bankruptcy	Rule	Behaviour	Bankruptcy
1	$(g-b)/(g+g)$	Never	16	$(g-b)/(g*g)$	$1=g>b$
2	$(g-b)/(g+b)$	$g>b=0$	17	$(g-b)/(g*b)$	$g>b=0$
3	$(g-b)/(g+k)$	Never	18	$(g-b)/(g*k)$	$g>b/(1-k)$
4	$(g-b)/(b+b)$	$g>3b$	19	$(g-b)/(b*b)$	$g>b^2+b$
5	$(g-b)/(b+k)$	$g>2b+k$	20	$(g-b)/(b*k)$	$g>b(1+k)$
6	$(g-b)/(k+k)$	$g>2k+b$	21	$(g-b)/(k*k)$	$g>b$
7	$(g-b)/(g-g)$	$g>b$	22	$(g-b)/(g/g)$	$g>b$
8	$(g-b)/(g-b)$	$g>b$	23	$(g-b)/(g/b)$	$g>b^2/(b-1)$
9	$(g-b)/(g-k)$	$g>b=0$	24	$(g-b)/(g/k)$	Never
10	$(g-b)/(b-g)$	Never	25	$(g-b)/(b/g)$	$g>b$
11	$(g-b)/(b-b)$	$g>b$	26	$(g-b)/(b/b)$	$g>b$
12	$(g-b)/(b-k)$	$g>2b-k$	27	$(g-b)/(b/k)$	$g>b(1+1/k)$
13	$(g-b)/(k-g)$	$b>g=0$	28	$(g-b)/(k/g)$	$g>b$
14	$(g-b)/(k-b)$	$g>b=0$	29	$(g-b)/(k/b)$	$g>b\geq 1$
15	$(g-b)/(k-k)$	$g>b$	30	$(g-b)/(k/k)$	$g>b$
31	$(g-b)/g$	$g>b=0$	33	$(g-b)/k$	$g>b$
32	$(g-b)/b$	$g>2b$	34	$(g-b)$	$g>b$

Table 1: Simple behaviour rules with conditions for positive risk of bankruptcy

matter how profitable it is do do so. The remaining 3 rules are exactly the ones that were represented by more than 50 per cent in our sample of most fit behaviour 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 favour of behaviour rules with these three features. However, this uncertainty proneness does impose a cost in terms of lost profits on the firms that host them, and in later periods they are replaced by firms with more rational behaviour rules. These firms not only behave in a manner which is consistent with expected utility maximization, they also act as if they were familiar with Bayesian statistics, despite the fact that they have no way of understanding the nature of the investment problem and no apparatus for estimating the relevant parameters of the unknown probability distributions. In the next section, we investigate whether this result carries over to a situation where the firms receive different amounts of information.

4.2 Case 2: Unqual information

We now consider the case where the firms still face the same success probability distribution in each period, but in contrast to the previous case, they no longer receive the same information about it. Recall that differences in information always imply differences in *amounts* of information, since two firms can only have different information if one of them has seen everything the other one has seen, plus some additional draws from the probability distribution p. In any case, it implies that different agents will typically have different opinions about the profitability of the risky investment alternative. This introduces an element of unsystematic risk, and one would therefore expect the firms to develop less risk averse behaviour in this case, as compared to the situation with equal information.[6]

To test this hypothesis, we ran another 20 experiments with the same GP–algorithm as before, except for drawing a seperate information state (g_t^i, b_t^i) for each firm, as described earlier. Again, we want to compare the evolved behaviour with the behaviour rule F^* defined in (10), which was shown to be the long–run survivor in the equal information case, as predicted by Bayesian decision theory. In the present case, however, there is no longer any apriori reason for believing that the behaviour rule F^* is still optimal, and we shall therefore refer to it as the *E-optimal* behaviour rule.

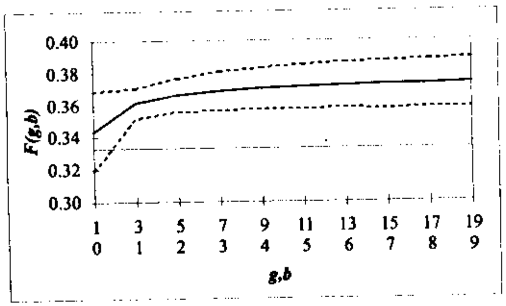

Figure 7: Mean behaviour over last 25.000 periods with 90 per cent confidence band for information states (g, b) such that $F^*(g, b) = 1/3$. (Unequal information)

Figure 7 displays the mean behaviour over the last 25.000 periods across all 20 runs, for our test set of information states (g, b) for which the E-optimal behaviour is 1/3. It is comparable to figure 5 of the previous subsection. As can be seen from the figure, the hypothesis that the mean behaviour is generated by the E-optimal behaviour rule can be

[6] Since expected growth is maximal for risk neutral investors, then if the population is infinite and there is only idiosyncratic risk, the population will be dominated by risk neutral investors in the long run. See Robson (1996) for an analysis of the finite population version of this result.

rejected. Except for the low information state $(g, b) = (1, 0)$, actual investment ratios are significantly higher than the E-optimal investment ratio of 1/3, as was to be expected from the introduction of asymmetric information.

Note however, that the firms tend to be less willing to invest in the risky alternative at low information states than at high ones. This is the opposite behaviour of what we got in the equal information case, hence it seems that the introduction of asymmetric information has caused a switch from uncertainty proneness to uncertainty aversion.

In order to check the robustness of this result, we did another 40 runs with the GP-algorithm. In 20 of them, we used a PopulationSize of 1000 and a BreedingRate of 2, and in the remaining 20, we used a PopulationSize of 4000 and a BreedingRate of 8. We then pooled all 60 runs and removed those 10 per cent that had the largest total deviation from E-optimal behaviour across the last 25.000 periods.

	r1000	r2000	r4000
Mean	0.358	0.361	0.362
Median	0.351	0.352	0.358
Standard Deviation	0.035	0.030	0.017
Minimum	0.306	0.329	0.344
Maximum	0.449	0.441	0.401
Count	19	16	19

Table 2: Summary statistics for 3 subsets of experimental data

Summary statistics for the three subsets of observations are given in table 2, based on the same data as figure 7. Each observation in each of the 3 data subsets consists of the mean behaviour for the population across the last 25.000 periods on average across the 10 cases in our test set of information states for which the E-optimal behaviour is 1/3. The main difference between them is that the sample variance decreases with increasing population size, which shows that the predictive power of the GP–algorithm improves with the size of the population.

Figure 8 shows the mean behaviour across all 54 runs on our test set of information states. As can be seen, the tendency to uncertainty aversion is no longer present in the extended data set. The mean investment ratio is almost constant across all 10 test cases, and hence consistent with Bayesian rationality. Moreover, the common investment ratio is approximately 10 per cent higher than the investment ratio of the E-optimal behaviour rule, which, as already mentioned, is consistent with the introduction of unsystematic risk.

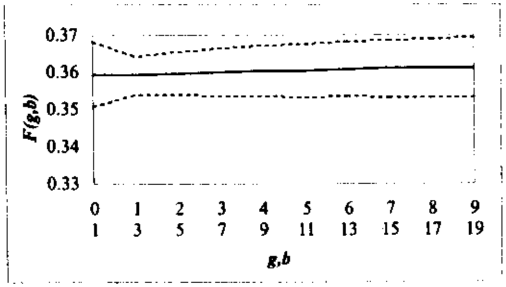

Figure 8: Mean behaviour over last 25.000 periods with 90 per cent confidence band for information states (g, b) such that $F^*(g, b) = 1/3$. (54 runs with unequal information)

4.3 Methodological issues

We believe these results show that Genetic programming can be a useful tool for studying the evolutionary basis for Bayesian rationality. 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 modelled as rigid rule followers who do not change their behaviour over time, except through bankruptcy or reorganization. As in the evolutionary game theory of Maynard Smith (1982), all change in behaviour takes place at the level of the population, but GP differs from evolutionary game theory by not requiring the modeller to specify all possible behaviour rules explicitly ex ante. In GP, new behaviour rules emerge as a result of random recombinations of behaviour 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 behaviour rules produced by GP are the result of learning by example, and generalization to the whole 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.

5 Conclusion

The purpose of this paper has been to investigate the usefulness of Genetic Programming as a tool for generating hypotheses about rational behaviour 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.

We have shown that when the agents have symmetric information, the algorithm systematically generates behaviour 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.

When we complicate matters by considering asymmetric information, the evolved behaviour becomes less risk averse, which is consistent with the unsystematic risk that is produced by the information asymmetry. Moreover, the behaviour is still consistent with Bayesian rationality, in the sense of being insensitive to changes in the amount of information that yield the same expected value for the unknown probability distributions.

Our interest in Knightean uncertainty is 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 behaviour in games, using the same technique that was shown to generate rational behaviour under Knightean uncertainty.

References

Aumann, R. and A. Brandenburger (1995), "Epistemic conditions for Nash equilibrium", *Econometrica 63*:1161–1180.

Alchian, A. A. (1950), "Uncertainty, evolution and economic theory", *Journal of Political Economy 58*:211–221.

Axelrod, R. (1987), "The evolution of strategies in the iterated prisoner's dilemma", In Davis og Lawrence (eds.) *Genetic Algorithms and Simulated Annealing* (Pittman)

Breiman, L. (1961), "Optimal gambling systems for favorable games", *Fourth Berkeley Symposium on Mathematical Statistics and Probability 1*:65–78.

Ellsberg, D. (1961) "Risk, ambiguity, and the Savage axioms", *Quarterly Journal of Economics 75*:643–669.

Friedman M. (1953), *Essays in Positive Economics* (Chicago: University of Chicago Press)

Gilboa, I, and D. Schmeidler (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* (Ann Arbor: University of Michigan Press)

Kinnear, K. E. (1994) (ed.) *Advances in Genetic Programming* (Cambridge: MIT Press)

Knight, F.H., (1921) *Risk, Uncertainty and Profit*, (Boston: Houghton Mifflin).

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

Koza, J. R. (1992) *Genetic Programming* (Cambridge: MIT Press)

Latané, H. A. (1959) "Criteria for choice among risky assets", *Journal of Political Economy*:144–155.

Marimon, R., E. McGrattan og T. J. Sargent (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: Cambridge University Press)

Nelson, R. R. og S. G. Winter (1982) *An Evolutionary Theory of Economic Change* (Cambridge: Belknap)

Robson, A.J. (1996), "A biological basis for expected and non-expected utility", *Journal of Economic Theory 68*: 397-424

Schumpeter, J. A. (1942), *Capitalism, Socialism and Democracy* (New York: Harper)

CES Working Paper Series

---

63 Tõnu Puu, The Chaotic Monopolist, August 1994

64 Tõnu Puu, The Chaotic Duopolists, August 1994

65 Hans-Werner Sinn, A Theory of the Welfare State, August 1994

66 Martin Beckmann, Optimal Gambling Strategies, September 1994

67 Hans-Werner Sinn, Schlingerkurs - Lohnpolitik und Investitionsförderung in den neuen Bundesländern, September 1994

68 Karlhans Sauernheimer and Jerome L. Stein, The Real Exchange Rates of Germany, September 1994

69 Giancarlo Gandolfo, Pier Carlo Padoan, Giuseppe De Arcangelis and Clifford R. Wymer, The Italian Continuous Time Model: Results of the Nonlinear Estimation, October 1994

70 Tommy Staahl Gabrielsen and Lars Sørgard, Vertical Restraints and Interbrand Competition, October 1994

71 Julia Darby and Jim Malley, Fiscal Policy and Consumption: New Evidence from the United States, October 1994

72 Maria E. Maher, Transaction Cost Economics and Contractual Relations, November 1994

73 Margaret E. Slade and Henry Thille, Hotelling Confronts CAPM: A Test of the Theory of Exhaustible Resources, November 1994

74 Lawrence H. Goulder, Environmental Taxation and the "Double Dividend": A Reader's Guide, November 1994

75 Geir B. Asheim, The Weitzman Foundation of NNP with Non-constant Interest Rates, December 1994

76 Roger Guesnerie, The Genealogy of Modern Theoretical Public Economics: From First Best to Second Best, December 1994

77 Trond E. Olsen and Gaute Torsvik, Limited Intertemporal Commitment and Job Design, December 1994

78 Santanu Roy, Theory of Dynamic Portfolio Choice for Survival under Uncertainty, July 1995

79 Richard J. Arnott and Ralph M. Braid, A Filtering Model with Steady-State Housing, April 1995

80 Vesa Kanniainen, Price Uncertainty and Investment Behavior of Corporate Management under Risk Aversion and Preference for Prudence, April 1995

81 George Bittlingmayer, Industry Investment and Regulation, April 1995

82 Richard A. Musgrave, Public Finance and Finanzwissenschaft Traditions Compared, April 1995

83 Christine Sauer and Joachim Scheide, Money, Interest Rate Spreads, and Economic Activity, May 1995

84 Jay Pil Choi, Preemptive R&D, Rent Dissipation and the "Leverage Theory", May 1995

85 Stergios Skaperdas and Constantinos Syropoulos, Competing for Claims to Property, July 1995

86 Charles Blackorby, Walter Bossert and David Donaldson, Intertemporal Population Ethics: Critical-Level Utilitarian Principles, July 1995

87 George Bittlingmayer, Output, Political Uncertainty, and Stock Market Fluctuations: Germany, 1890-1940, September 1995

88 Michaela Erbenová and Steinar Vagstad, Information Rent and the Holdup Problem: Is Private Information Prior to Investment Valuable?, September 1995

89 Dan Kovenock and Gordon M. Phillips, Capital Structure and Product Market Behavior: An Examination of Plant Exit and Investment Decisions, October 1995

90 Michael R. Baye, Dan Kovenock and Casper de Vries, The All-pay Auction with Complete Information, October 1995

91 Erkki Koskela and Pasi Holm, Tax Progression, Structure of Labour Taxation and Employment, November 1995

92 Erkki Koskela and Rune Stenbacka, Does Competition Make Loan Markets More Fragile?, November 1995

93 Koji Okuguchi, Effects of Tariff on International Mixed Duopoly with Several Markets, November 1995

94 Rolf Färe, Shawna Grosskopf and Pontus Roos, The Malmquist Total Factor Productivity Index: Some Remarks, November 1995

95 Guttorm Schjelderup and Lars Sørgard, The Multinational Firm, Transfer Pricing and the Nature of Competition, November 1995

96 Guttorm Schjelderup, Kåre P. Hagen and Petter Osmundsen, Internationally Mobile Firms and Tax Policy, November 1995

97 Makoto Tawada and Shigemi Yabuuchi, Trade and Gains from Trade between Profit-Maximizing and Labour-Managed Countries with Imperfect Competition, December 1995

98 Makoto Tawada and Koji Shimomura, On the Heckscher-Ohlin Analysis and the Gains from Trade with Profit-Maximizing and Labour-Managed Firms, December 1995

99 Bruno S. Frey, Institutional Economics: What Future Course?, December 1995

100 Jean H. P. Paelinck, Four Studies in Theoretical Spatial Economics, December 1995

101 Gerhard O. Orosel and Ronnie Schöb, Internalizing Externalities in Second-Best Tax Systems, December 1995

102 Hans-Werner Sinn, Social Insurance, Incentives and Risk Taking, January 1996

103 Hans-Werner Sinn, The Subsidiarity Principle and Market Failure in Systems Competition, January 1996

104 Uri Ben-Zion, Shmuel Hauser and Offer Lieberman, A Characterization of the Price Behaviour of International Dual Stocks: An Error Correction Approach, March 1996

105 Louis N. Christofides, Thanasis Stengos and Robert Swidinsky, On the Calculation of Marginal Effects in the Bivariate Probit Model, March 1996

106 Erkki Koskela and Ronnie Schöb, Alleviating Unemployment: The Case for Green Tax Reforms, April 1996

107 Vidar Christiansen, Green Taxes: A Note on the Double Dividend and the Optimum Tax Rate, May 1996

108 David G. Blanchflower and Richard B. Freeman, Growing Into Work, May 1996

109 Seppo Honkapohja and George W. Evans, Economic Dynamics with Learning: New Stability Results, May 1996

110 Seppo Honkapohja and George W. Evans, Convergence of Learning Algorithms without a Projection Facility, May 1996

111 Assar Lindbeck, Incentives in the Welfare-State, May 1996

112 Andrea Ichino, Aldo Rustichini and Daniele Checchi, More Equal but Less Mobile?, June 1996

113 David Laidler, American Macroeconomics between World War I and the Depression, June 1996

114 Ngo Van Long and John M. Hartwick, Constant Consumption and the Economic Depreciation of Natural Capital: The Non-Autonomous Case, August 1996

115 Wolfgang Mayer, Gains from Restricted Openings of Trade, August 1996

116 Casper de Vries and Jón Daníelsson, Tail Index and Quantile Estimation with Very High Frequency Data, August 1996

117 Hans-Werner Sinn, International Implications of German Unification, October 1996

118 David F. Bradford, Fixing Capital Gains: Symmetry, Consistency and Correctness in the Taxation of Financial Instruments, October 1996

119 Mark Hallerberg and Scott Basinger, Why Did All but Two OECD Countries Initiate Tax Reform from 1986 to 1990?, November 1996

120 John Livernois and C. J. McKenna, Truth or Consequences? Enforcing Pollution Standards, November 1996

121 Helmut Frisch and Franz X. Hof, The Algebra of Government Debt, December 1996

122 Assaf Razin and Efraim Sadka, Tax Burden and Migration: A Political Economy Perspective, December 1996

123 Torben M. Andersen, Incomplete Capital Markets, Wage Formation and Stabilization Policy, December 1996

124 Erkki Koskela and Rune Stenbacka, Market Structure in Banking and Debt-Financed Project Risks, December 1996

125 John Douglas Wilson and Xiwen Fan, Tax Evasion and the Optimal Tax Treatment of Foreign-Source Income, January 1997

126 Terje Lensberg, Investment Behaviour under Knightian Uncertainty - An Evolutionary Approach, January 1997