

A genetic algorithm approach to farm investment[†]

Oscar Cacho and Phil Simmons*

A farm portfolio model is specified with two risky enterprises and a risk-free asset which may be held short or long by the farmer. The model is solved numerically using a genetic algorithm. It is shown that the assumption of competitive adaptation leads to a violation of normative efficiency. Those who survive are not the most efficient in a normative sense.

1. Introduction

Genetic algorithm (GA) models have been used in a number of disciplines in the last two decades with applications in physics, electrical and chemical engineering, and the biological and behavioural sciences including economics. In general, they are used for two purposes. First, as a methodological tool to maximise functions that are either highly non-linear or very large in terms of number of control variables. Second, GA models are used in the social sciences to test specific behavioural hypotheses. In economics, GA models have been used in both these contexts with applications in monetary theory, index design, dynamic cobweb and stock market behaviour, optimisation under imperfect competition, and as learning models (Arifovic 1994; Arifovic 1995; Farley and Jones 1994; Holland and Miller 1991; Marks 1992). In this article, a GA model is used to examine farm investment.

Szpiro (1997) examined investment behaviour using a GA model of a stock market and found that when firms were selected against on the basis of poor profit performance, they developed caution in their investment strategies. The results are interesting since they reflect an alternative approach to risk behaviour based on competitive adaptation rather than maximising certainty equivalent value. In our study, we use a similar

[†] The authors wish to thank John Dillon and an anonymous referee for helpful comments on earlier drafts of this article.

*Oscar Cacho and Phil Simmons both work at the Department of Agricultural and Resource Economics, University of New England, Armidale, NSW 2351, Australia.

approach to examine farm investment. The goals in our study were, first, to adapt Szpiro's model to resemble a farm rather than a stock market and, second, to augment the model so that a test for normative efficiency could be undertaken. The purpose of this article is not to present an alternative optimisation technique for farm managers, but rather to raise important questions concerning the risk attitudes that are likely to evolve under alternative institutional arrangements, as they affect the risk exposure of producers.

In the following section we develop a simple economic model consisting of agents who borrow and use their capital to produce a mix of two crops. The implementation of the model within a genetic algorithm (GA) framework is described in section 3. A set of experiments designed to study a number of questions is described in section 4. The behaviour of the GA model is analysed in detail, and the results of the optimisation (or evolutionary) process are presented in sections 5 and 6, the later section also containing a test of the Separation Theorem. The final section presents discussion and conclusions and identifies some unresolved issues.

2. Farm model and economic environment

A portfolio investment model was specified with constant returns to scale, two crops, s and w , and the possibility of either borrowing or lending. The expected return from farm capital is r_a :

$$r_a = p_s r_s + p_w r_w \quad (1)$$

where r_s and r_w are expected returns from s and w , respectively and p_s and p_w are proportions of farm capital allocated to s and w . It is assumed p_s and p_w are non-negative, r_s and r_w are independent so $\text{Cov}(r_s, r_w) = 0$ and there is an additivity restriction, $p_w = 1 - p_s$.

The farmer borrows or lends at 'risk-free' rate r_b so the expected return to equity, r_e , is:

$$r_e = p_a r_a + p_b r_b \quad (2)$$

where p_a is the ratio of farm value to equity and p_b is the ratio of borrowing or lending to equity. Hence, $0 \leq p_b \leq 1$ if the farmer is a lender and $p_b < 0$ if the farmer is a borrower. Thus, p_a is non-negative to prevent 'shorting' of farm capital and the additivity restriction is $p_a = 1 - p_b$. With limited borrowing $-b_{\max} \leq p_b \leq 1$, where b_{\max} is the borrowing limit.

Farm income is defined as return to equity by assuming one unit of equity capital and, after substituting the two additivity restrictions and equation 1 into equation 2, r_e becomes:

$$r_e = p_b r_b + (1 - p_b)(p_s r_s + (1 - p_s)r_w). \quad (3)$$

The farmer chooses both the crop mix and the amount of borrowing or lending, hence the decision variables are p_b and p_s . Returns to s and w are independent and subject to separate market and production shocks, u_{1j} and u_{2j} ($j = s$ or w), respectively:

$$r_j = e^{u_{1j}+u_{2j}} - c_j \quad (4)$$

where $u_{ij} \sim N(\mu_{ij}, \sigma_{ij})$ ($i = 1, 2; j = s, w$) and $c_j > 0$ is a constant. Hence, returns from both crops have three-coefficient, bivariable lognormal distributions bounded below by c_j and unbounded above. It is further assumed production risk is correlated across farms and that farmers, collectively, are price makers with $\text{Corr}(u_{1j}, u_{2j}) = \rho_{12j} \leq 0$ (Anderson, Dillon and Hardaker 1977, pp. 171–2).

3. The genetic algorithm

There are many different types of GA models. Their defining characteristic is that they contain the elements of evaluation, selection, crossover and mutation (Goldberg 1989). These elements are explained below in terms of the farm model.

The numerical model consists of a population of 100 farmers (agents) with identical technology and facing the same risks. Each agent is assigned starting values for two ‘genes’, p_b and p_s , which are drawn randomly from a uniform distribution with bounds $-b_{\max} \leq p_b \leq 1$ and $0 \leq p_s \leq 1$, where b_{\max} , the maximum debt allowed, was assigned a value of 5. The starting values for p_s are consistent with the non-negativity constraint preventing shorting of farm capital and the additivity restriction, $p_s = 1 - p_w$. Each gene is represented by a ten-character binary string, which maps into the set of integers $1 - 2^{10}$. The real-valued parameters p_b and p_s are mapped into this discrete space, which contains 1024 points, resulting in sampling intervals of 5.86×10^{-3} for p_b and 9.77×10^{-4} for p_s .

The agents undertake borrowing or lending and invest by allocating capital between s and w over k periods according to their assigned starting values for p_b and p_s . Investment occurs in a stochastic environment resulting from the distributions of r_s and r_w as specified in equation 4. Different values of r_j ($j = s, w$) are obtained for each period based on u_{1j} and u_{2j} which, in turn, are drawn from a bivariate normal distribution with correlation ρ_{12j} .

At the end of the k -th period, the investment performance of each agent is evaluated using a value function calculated by summing return to equity over the investment horizon:

$$R_i = \sum_1^k r_{eik} \quad (5)$$

where r_{eik} for agent i in period k is estimated by equation 3. A fitness function is then calculated by subjecting the value function (5) to linear scaling as described by Goldberg (1989, pp. 76–9), so that the agent with the highest ‘raw fitness’ (R) score produces an expected two offspring, whereas the ‘average’ agent produces one expected offspring per generation. Scaling may prevent two problems: premature convergence early in a run, and slow convergence in later generations. Agents are then selected for reproduction using classical roulette wheel selection, where the probability of selection is proportional to the fitness of the individual relative to the rest of the population (Mitchell 1997, pp. 166–7). Genes belonging to agents who are not selected for reproduction disappear from the population.

After selection, crossing over occurs with each pair producing two offspring and then disappearing, leaving population size in the second generation the same as in the first. Crossover allows transmission of genes from one generation to the next and facilitates evolution of agents better adapted to their stochastic environment. The crossover operation makes copies from the genes of the two parents using the ‘bit string swapping’ mechanism described in Goldberg (1989) and elsewhere. The probability of single-point crossover occurring was set at 0.6, which means that there was a 0.4 probability of offspring being identical to their parents.

Unlike in Szpiro (1997), there is no inheritance of wealth. Only expertise is inherited and each generation starts out with the same technical possibilities and level of equity capital.

Pairing and crossover ensure two important things happen in the GA model. The first is, based on the value function (5), poor performers are removed from the population so, eventually, only the fittest survive. The second is, because pairing involves the whole population, inferior genes can survive, albeit in proportions decreasing with each generation, and hence the model has a ‘genetic memory’. If the environment changes then an inferior gene may cease to be inferior and become dominant within a few generations. However, a gene eventually disappears if it does not contribute to overall population fitness.

A small proportion of the new generation may undergo random mutation of one or more of their genes. Mutation consists of flipping a random bit (i.e. a zero changes to a one and vice versa) in the binary representation of the parameter, with the probability of a bit flipping in our model being 0.01. As this evolutionary process moves through time, the genetic make-up of the population converges to values of p_b and p_s which provide the best survival value for the given parameters r_b , μ_{ij} , σ_{ij} , ρ_{12j} , c_j and k .

4. Experimental design

A series of experiments were run, each consisting of 200 generations with a population of 100 agents. The base case was designed to determine whether the stochastic environment would cause agents to evolve risk-averse behaviour, since one of the objectives of the study was to examine competitive adaptation under different assumptions about risk. There is a number of ways to adjust risk in the GA framework. One approach is to increase the period of evaluation (investment horizon k) which reduces the variance of the value function. A second approach is to vary the severity of the selection process by adjusting the probability of survival. These variations were tested in the experiments described below.

4.1 Experiment 1: standard environment

Values for fourteen coefficients were chosen for the experiments as shown in table 1. This meant that r_s and r_w had expected values of 9.227 and 8.684 and standard deviations of 7.486 and 4.722, respectively. Treatments within each experiment consisted of values of the risk-free rate (r_b) ranging from 8 to 10. This range contains the critical value at which a risk-neutral individual would switch from borrowing to lending, this critical value (r_{crit}) is equal to 9.227, the expected return of the most profitable crop. This scenario was taken to represent the 'standard' environment and is denoted by S.

4.2 Experiment 2: harsh environment

The severity of selection was increased by reducing the fitness value of agents who had negative wealth at the end of the investment horizon (this could represent an unforgiving creditor). This was implemented as an arbitrary 'debt penalty' which was imposed on individuals who were in debt at the end

Table 1 Parameter values used in experiments

Parameter	Value	Description
k	3	investment horizon
b_{max}	5	borrowing limit
r_b	8–10	risk-free rate
$\mu_{1w} = \mu_{2w}$	1.28	mean of price and yield shocks for crop w
$\sigma_{1w} = \sigma_{2w}$	0.25	standard deviation of price and yield shocks for crop w
c_w	5	cost of producing crop w
$\mu_{1s} = \mu_{2s}$	1.38	mean of price and yield shocks for crop s
$\sigma_{1s} = \sigma_{2s}$	0.31	standard deviation of price and yield shocks for crop s
c_s	8	cost of producing crop s
$\rho_{12w} = \rho_{12s}$	−0.1	price-yield correlation

of the investment horizon. The penalty consisted of cancelling out the raw fitness score of the worst-performing individual (thereby preventing its reproduction) and interpolating linearly for the remaining debtors. Specifically:

$$d_i = (R_i < 0) \cdot |R_i| \quad (6)$$

where d_i is the penalty, the term in brackets assumes a value of one if true and zero if false, a scaling factor (w) is then defined as:

$$w_i = 1 - \frac{d_i}{\max(d)} \quad (7)$$

and the fitness function is:

$$F_i = [\beta_0 + \beta_1 R_i] \cdot w_i \quad (8)$$

where the β s are linear scaling parameters estimated as in Goldberg (1989) to obtain the desired number of offspring from the fittest individual. This scenario was taken to represent a harsh environment and is denoted by H.

4.3 Experiment 3: mild environment

Increasing the investment horizon (k) is equivalent to reducing hostility in the environment and is hypothesised to make farmers less cautious. Hence, in the third experiment, k was increased from three to 15 periods. This allowed agents to offset losses in poor periods against gains from favourable periods prior to evaluation and selection. This experiment represents a mild environment and hereafter is denoted by M.

5. Model behaviour

The GA model was run under the assumptions described above and average values for p_b and p_s were obtained. The model was first run in a deterministic environment to ensure that it would behave as predicted by theory. The model was then run in stochastic mode. The results of these tests are described below.

5.1 Deterministic environment

The model was solved for three values of r_b (8, 9.3 and 10) in a deterministic environment, simulated by setting all σ_{ij} to zero and adjusting the μ_{ij} parameters to yield the same expected values as with the stochastic model described in table 1 ($\mu_{1s} = \mu_{2s} = 1.423$ and $\mu_{1w} = \mu_{2w} = 1.308$). The GA converged rapidly to corner solutions (figures 1A and 1C) for p_b and p_s . For all three values of r_b , convergence was accompanied by a rapid drop in

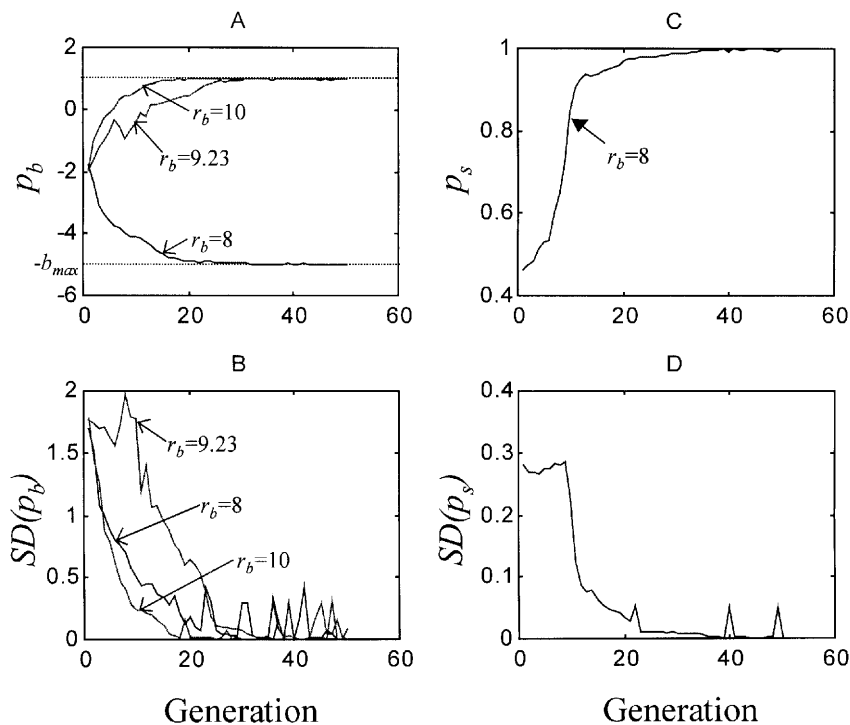


Figure 1 Results from deterministic runs at three values of the risk-free rate

genetic variability as represented by the standard deviations (SD) of p_b and p_s (figures 1B and 1D). This was achieved as the genes of the fittest individuals took over an increasing proportion of the population. Convergence occurred within 16 to 30 generations depending on the value of r_b (figure 1). Convergence was slower when r_b was close to the critical value of 9.227, a result of selection pressure on specific genes being weaker. By the 30th generation, variability had decreased to negligible values, with a small amount of noise remaining from ongoing mutation.

Results were as expected. With the risk-free rate below the farm return ($r_b = 8$), agents borrow as much as possible (figure 1A) and invest all available capital on the crop with the highest expected return (figure 1C). When $r_b > E(r_s)$, no production occurs and agents lend all their capital and, under these conditions, the value of p_s is irrelevant.

5.2 Stochastic environment

The concept of the fittest individual becomes somewhat blurred within a stochastic environment. Whereas in the deterministic case individuals with

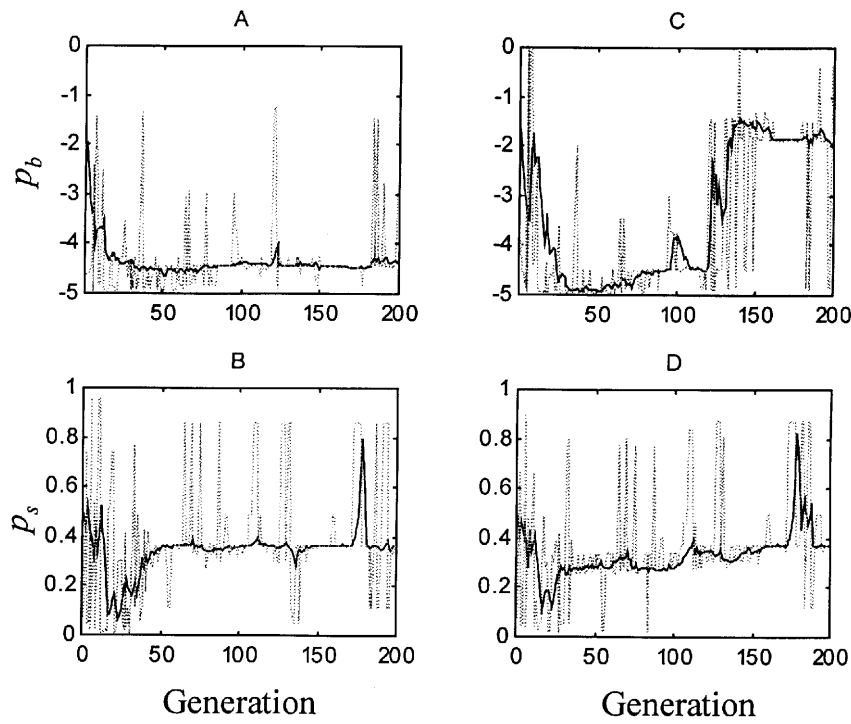


Figure 2 Behaviour of GA under stochastic conditions with $r_b = 8$

Note: A and B standard environment, C and D harsh environment. Dotted lines represent the best individuals of each generation and solid lines represent the average of the population.

the ‘right’ parameter values always perform better than the population average, in the stochastic case a set of genes which performs well in one period may perform poorly the next. ‘Performance’ depends on prevailing and not just average prices and yields. This is illustrated in figure 2 where the ‘best’ individual (the one which accumulated the most wealth) in any given generation varies sharply through time.

A convergence of sorts occurs in the stochastic model and individuals which have evolved over 200 generations exhibit risk-averse behaviour in the sense that they borrow less than the maximum allowed and invest only part of their capital on the high-risk, high-return crop (figure 2). In the standard environment (S) convergence occurred at $p_b = -4.34$ and $p_s = 0.36$ which implies that $p_w = 0.64$ (figures 2A and 2B). In the harsh environment (H), convergence occurred at $p_b = -1.82$ and $p_s = 0.37$ (figures 2C and 2D). Thus individuals that evolved under harsher conditions borrowed less and might be viewed as more risk averse. The similar values of p_s indicate that as the

environment became harsher, agents responded by adjusting borrowing rather than their cropping mix. This is relevant to the Separation Theorem which is discussed later.

5.3 Convergence and the role of mutation

Convergence in a GA can be declared when the entire population has evolved to the same genetic make-up, within the desired tolerance. Strictly speaking, convergence never occurs in the GA described here; the presence of a stochastic environment means that genes that are dominated in a given generation may be dominant in the next; and the probability of mutation means that new genes may emerge at any time. An illustration of the role of mutation in preventing early convergence is presented in figure 3. With $r_b = 9.23$ the GA had practically converged to -2.0 by the 70th generation,




Figure 3 Results with $r_b = 9.23$ in standard environment

Note: Figure A shows the mean value of p_b for a population of 100 individuals, figure B shows the standard deviation.




Figure 4 Spread of surviving agents in parameter space under alternative assumptions

Note: *a*: $r_b = 8$, deterministic; *b*: $r_b = 10$, deterministic; *c*: $r_b = 8$ standard environment, stochastic; *d*: $r_b = 8$, harsh environment, stochastic. Each marked area contains 100 points.

until a large mutation occurred in the 120th generation and redirected the algorithm to converge to the true optimum. The mutation caused a large increase in the variability of p_b (figure 3B), and a rapid change in the average values of p_b (from -2 to 1.0) over a period of 15 generations (figure 3A).

The fact that the initial population is drawn from a uniform distribution has two effects: initial average parameter values occur close to the centre of the feasible zone, and initial variability among individual agents is high (see figure 1). A plot of the initial population in parameter space would have points randomly scattered throughout the feasible area. But, after being subject to selection pressure over 200 generations, the parameter values of surviving agents are concentrated within a small area in the feasible space (figure 4), whose location depends on the environment in which the agents evolved. The movement away from risk-neutrality is illustrated by the shift, upwards and to the left, from area *a* to *c* and *d* (figure 4). Area *b* is shown simply to illustrate that convergence was reached in the lending sense but not in the crop mix sense when $r_b > E(r_s)$; this is because no production occurs.

The results presented in this section provide insight into the behaviour of the GA, but they provide no conclusive evidence. These results represent the outcome obtained from a single numerical experiment, one element within an infinite possible set of price and yield realisations. The analysis is now extended to include replication in the context of Monte Carlo experiments.

6. Monte Carlo experiments

Monte Carlo experiments were performed by solving the GA in a 4×3 factorial design. Four values of r_b (8, 9.2, 9.3 and 10) were tested in each of the three environments (M , S , H). Treatments are denoted by their environment- r_b combination (i.e. M -8 or S -9.2). One hundred replicates were obtained for each treatment by running the model with a different random seed each time. This design resulted in 12 different treatments, each yielding 10 000 surviving individuals (100 individuals per experiment \times 100 experiments).

As expected, risk aversion increased as the environment changed from mild to harsh. With $r_b = 8$ the average value of p_b increased from -4.97 to -4.28 to -2.67 as the environment changed from mild to standard to harsh (figure 5 and table 2). A similar pattern was obtained with $r_b = 9.2$, but at




Figure 5 Effect of the risk-free rate (r_b) on the average borrowing behaviour of surviving agents under mild (M), standard (S) and harsh (H) environments

Table 2 Average results of Monte Carlo experiments ($n = 10\,000$)

	Environment		
	Mild	Standard	Harsh
With $r_b = 8.0$			
mean p_b	-4.97	-4.28	-2.67
mean p_s	0.76	0.55	0.46
mean R	14.04	14.13	11.74
risk aversion index	0.03	0.72	2.33
With $r_b = 9.2$			
mean p_b	-0.18	-0.13	0.41
mean p_s	0.83	0.61	0.64
mean R	8.85	8.80	9.09

higher parameter values, mean p_b values increased from -0.18 to -0.13 and then to 0.41 as the environment changed from mild to standard to harsh (figure 5 and table 2). Thus, the harsh environment led to the evolution of an 'average' agent lending 41 per cent of its capital when the risk-free rate is slightly below the critical value, whereas a classically risk-neutral individual still borrows as much as possible under these circumstances. Agents that evolved in the standard and mild environments also exhibited risk-averse behaviour, borrowing less than 4 per cent of the maximum allowed (figure 5).

Risk aversion also evolved in the crop-mix choice. Mean p_s values at $r_b = 8$ decreased from 0.76 to 0.55 and then to 0.46 as the environments changed from mild to standard to harsh (table 2).

Interestingly, as the value of the risk-free rate increased slightly above the critical level ($r_b = 9.3$), a mild gambling behaviour evolved in all three environments. Only between 0.19 and 0.49 of capital available was lent at the (higher) risk free rate, the remainder being invested in farm production which yields a lower expected return. This suggests that the compensation provided by occasional high prices and/or yields makes the gamble attractive. This mild risk-preferring behaviour did not occur at the higher r_b value of 10, where all agents lent all their capital.

Similarly to Szpiro (1997), the distance between the maximum borrowing rate and the actual borrowing rate adopted by the surviving population can be taken as an arbitrary measure of risk aversion. By this measure, the level of risk aversion ranges from 0.03 in the mild environment to 2.33 in the harsh environment (table 2).

So far we have described only average results for a set of twelve populations of 10 000 surviving individuals each, without discussing other

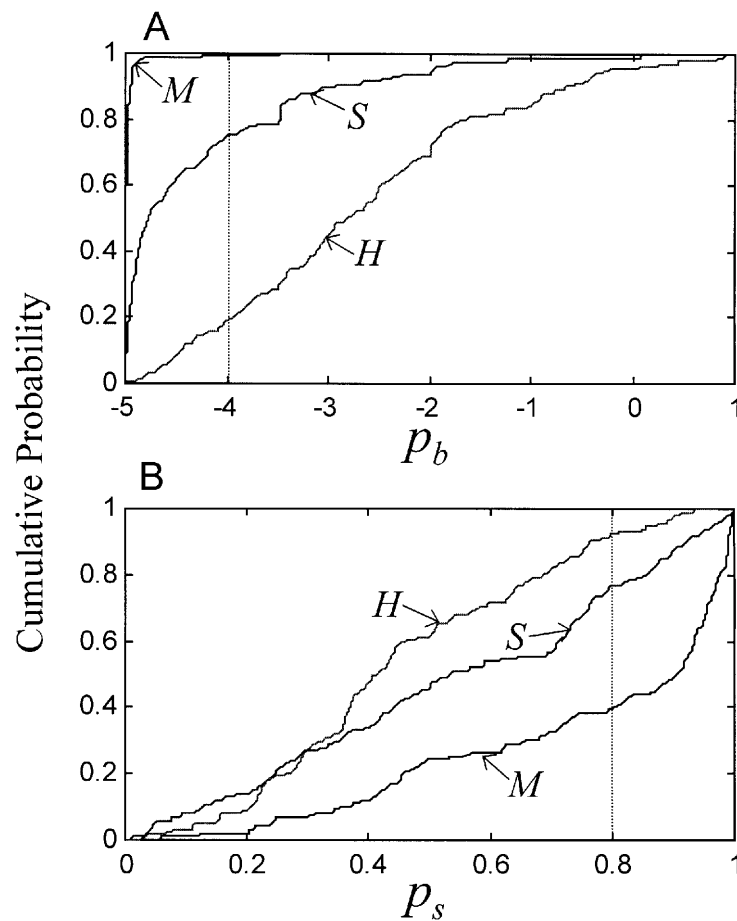


Figure 6 Cumulative density functions of parameter values with $r_b = 0.8$

moments of the distributions. Inspection of cumulative density functions reveals that there is considerable variation among surviving individuals in each population (figure 6). However, clear patterns still emerge in support of the previous discussion. The proportion of surviving agents who borrowed more than 4 times the available equity ($p_b < -4$) was practically 1.0 in the mild environment, decreasing to 0.75 in the standard environment and 0.2 in the harsh environment (figure 6A). In the same vein, the proportion of surviving agents who invested more than 0.8 capital in the high-return, high-risk crop was 0.55 in the mild environment, decreasing to 0.2 and 0.05 in the standard and harsh environments (figure 6B).




Figure 7 Graphical representation of the Separation Theorem

6.1 The Separation Theorem

The Separation Theorem, attributed to Tobin, states, in the context of farming, that if capital markets are efficient, and there exists a risk-free asset, then the crop mix is not influenced by risk preferences. This is illustrated in figure 7, where SW is the farm income-risk frontier and $r_b z$ is the risk-efficient frontier comprising different combinations of risky farm capital and risk-free asset, the latter held short or long. An efficient farmer is restricted to the risk-efficient frontier as indicated by the indifference curves and hence risk attitudes are separated from the crop-mix decision. We have seen that the final distribution of p_s differs among the three environments. Thus, an interesting question is whether the Separation Theorem holds in our model. The results presented above cannot be used to test this hypothesis, because each of the three environments resulted in a different risk-efficient frontier.

To test the Separation Theorem we must use a given risk-efficient frontier and allow only the level of risk aversion to vary. Thus the fitness function must be redefined to account for risk aversion, and a functional form must be assigned to solve the numerical model. A convenient form for the utility function is (Anderson, Dillon and Hardaker 1997, p. 99):

$$U_i = 1 - \exp(-\alpha R_i) \quad (9)$$

where α is a coefficient of risk aversion. The hypothesis that the Separation Theorem holds was tested by using this function to measure fitness and solving Monte Carlo experiments for two different values of α (0.05 and 0.16) and with r_b at 8. The spread in α values used in these experiments is

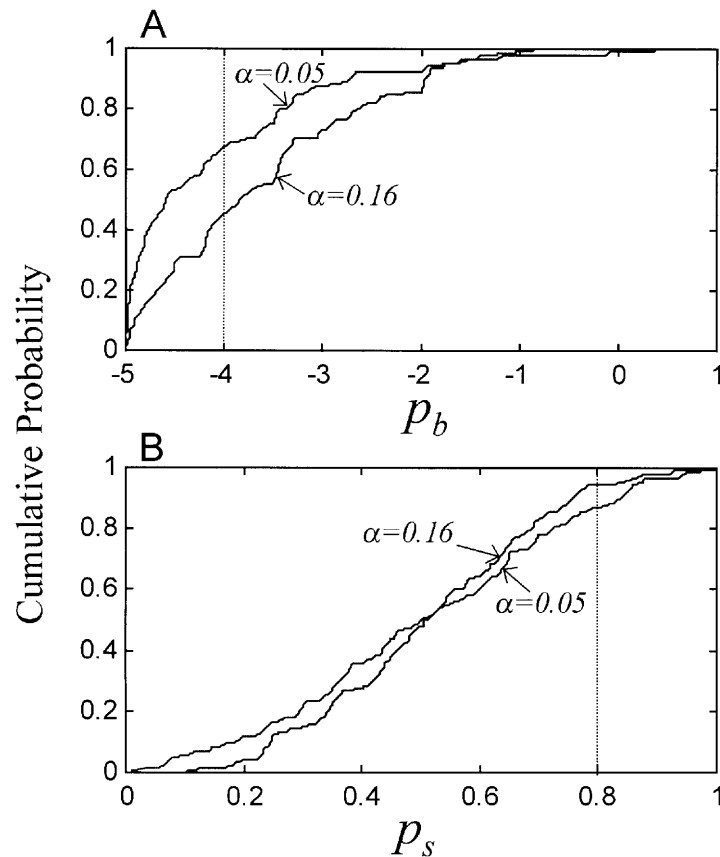


Figure 8 Cumulative density functions of parameter values under utility maximisation (with $r_b = 0.8$)

Note: Lines represent two different levels of risk aversion coefficient α in equation 9.

proportional to the spread in risk aversion indexes calculated between the standard and harsh environments (see table 2).

As expected, higher risk aversion resulted in lower borrowings (figure 8A). The proportion of the surviving population that borrowed within one unit of the limit ($p_s \leq -4$) decreased from 0.67 to 0.46 per cent as α increased from 0.05 to 0.16. Results on crop mix are not as clear-cut as those on borrowing (figure 8B). The density functions of the two experiments cross and their means are similar (0.507 and 0.515 at α of 0.05 and 0.16); however, the hypothesis that these means are equal was rejected ($p < 0.05$) by an F-test ($F = 7.26$). Note that the proportion of surviving individuals who invested less than 0.8 of their capital in the high-risk crop increased from 0.87 to 0.95 as α increased from 0.05 to 0.16 (figure 8B).

7. Conclusion

The cautious behaviour exhibited by the agents which evolved in our experiments comes directly from selection. If farm borrowings are high then a few poor market or production periods lead to low or negative incomes so selection favours neighbours who borrowed less. Alternatively, if borrowings are too low, returns are not competitive in normal and good periods. The fittest agents were those who struck a balance between too much and too little caution. These results were qualitatively similar to those that come from maximising certainty equivalent income with risk aversion.

An increase in the investment horizon caused agents' attitudes to move closer to risk neutrality, whereas increasing the harshness of the environment decreased both the amount borrowed and the proportion invested in the high-return high-risk crop.

The inclusion of a second risky crop in the model allowed a test for normative efficiency based on the Separation Theorem. The change in crop mix, p_s , occurring in the experiments with changes in the selection pressure do not mean the Separation Theorem was violated, because each environment resulted in a different possibility frontier. However, a test based on an arbitrary utility function and using two plausible risk-aversion coefficients, suggests that the Separation Theorem may have been violated. Hence, our results suggest that agents who are competitively adaptive may not be normatively efficient. This implies that the most fit agents may not be the most efficient and the most efficient may not be the most likely to survive.

Mutation is important in genetic algorithms designed for function maximisation, because it prevents the population from converging to a local maximum. Strictly speaking, the model developed here belongs to this class, since the fitness function leads to wealth maximisation. In the context of the stochastic environment prevailing in our model, mutation has the role of allowing agents to explore strategies which might otherwise have been left untested. The downside of mutation is that true convergence is never achieved. To prevent this problem, non-uniform mutation operators have been developed (Michalewicz 1994). These operators improve the fine-tuning capabilities of the GA by decreasing the impact of mutation as the final generation approaches. This scheme encourages a thorough search of 'genetic space' in early generations while constraining further search to local areas in late generations. This mutation strategy was not used in our model, but it may prove useful in future work.

The major problem with a GA approach to investment is that its theoretical basis, while powerful, is not sufficiently developed for the types of applied problems agricultural economists are interested in. However, the GA model does have some conceptual features that are attractive. The

genetic memory is an improvement on the treatment of priors in certainty equivalent models where 'memory' is usually restricted to perfect prior information or, at best, Bayesian or first or second-order ARCH processes. Another attractive feature of the GA risk model is that the researcher can model risk responses over time. While expected utility theorists have made progress here, most dynamic risk models are still solved as a string of single-period problems.

The most interesting question arising from the study is whether reconciliation of survival models and expected utility models is possible. There are a number of reasons why it is unlikely one theory will be found to be a special case of the other. First, the competitive adaptive model is based on 'survival of the fittest' while expected utility theory is based on satiation (diminishing marginal utility). These aspects of behaviour are very different from each other. Second, since fitness may not correspond to normative efficiency, any optimising agents introduced to the model, unless immune from selection, may disappear within a few generations. Basu (1996, p. 745), concluded from an evolutionary game-theory framework: 'fitness maximization is not a preference that will always emerge in equilibrium. Hence, at least for the time being, the attempt to bring utility and fitness into alignment must be abandoned as futile.' Chavas (1991) argues that optimising and non-optimising agents may co-exist in a non-identical population. However, from a GA perspective, it is not clear how his optimising agents derive their immunity when non-optimal agents are dominant. For a unifying framework to be developed, a central question is how farmers might trade likelihood of survival for income.

References

- Anderson, J.R., Dillon J.L. and Hardaker, B. 1977, *Agricultural Decision Analysis*, Iowa State University Press, Ames.
- Arifovic, J. 1994, 'Genetic algorithm learning and the cobweb model', *Journal of Economic Dynamics and Control*, vol. 18, no. 1, pp. 3–28.
- Arifovic, J. 1995, 'Genetic algorithms and inflationary economics', *Journal of Monetary Economics*, vol. 36, no. 1, pp. 219–43.
- Basu, K. 1996, 'Notes on evolution, rationality and norms', *Journal of Institutional and Theoretical Economics*, vol. 152, no. 4, pp. 739–49.
- Chavas, J-P., 1991, 'On sustainability and the economics of survival', *American Journal of Agricultural Economics*, vol. 75, no. 1, pp. 72–83.
- Farley, A.M. and Jones, S. 1994, 'A genetic algorithm to determine an index of leading indicators', *Computational Economics*, vol. 7, no. 3, pp. 163–73.
- Goldberg, D.E. 1989, *Genetic Algorithms in Search, Optimization, and Machine Learning*, Addison-Wesley, New York.
- Holland, J.H. and Miller, J.H. 1991, 'Artificial adaptive agents in economic theory', *American Economic Review: Papers and Proceedings of the 103rd Annual Meeting of the American Economic Association*, pp. 365–70.

- Marks, R.E. 1992, 'Breeding optimal strategies: optimal behaviour for oligopolists', *Journal of Evolutionary Economics*, vol. 2, no. 1, pp. 17–38.
- Michalewicz, Z. 1994, *Genetic Algorithms + Data Structures = Evolution Programs*, 2nd edn, Springer-Verlag, Berlin.
- Mitchell, M. 1997, *An Introduction to Genetic Algorithms*, MIT Press, Cambridge, Massachusetts.
- Szpiro, G. 1997, 'The emergence of risk aversion', *Complexity*, vol. 2, no. 4, pp. 31–9.