Optimisation of a Whole-farm Model

Neal, M.[1], Drynan, R., Fulkerson, W., Levy, G., Wastney, M., Post, E., Thorrold, B., Palliser, C., Beukes, P., and Folkers, C.

Keywords: Optimisation, Dairy, Model, Genetic Algorithm

The central aim of this project is to improve dairy farmer's strategic (long term) decisions. The process involved the integration of an economics component into an existing whole-farm model, followed by utility based optimisation using a genetic algorithm. The optimisation is first conducted with a single objective function assuming perfect capital and labour markets for a farmer with two investment alternatives: a dairy farm or risk-free asset. The problem is then broadened to a stochastic dominance problem, where the assumption of perfect capital markets is relaxed. Emphasis is then placed on post-optimisation analysis, especially through visual approaches.

1. Decision-making for dairy farmers

A dairy farmer must make both strategic and tactical (short term) decisions. Strategic decisions include choice of forage species, stocking rate, choice of animal genetics, calving pattern, expected quantity of purchased feeds and the share of capital to devote to plant and machinery. Tactical decisions could include choosing which paddock to feed on a particular day and with how much supplementary feed, what speed to rotate paddocks and how much fertiliser to apply over the next week. The particular focus of this paper is assisting the strategic decisions that the farmer must make, especially in reference to New Zealand dairy farms.

The farmer's decision problem is to choose a strategy, given their preferences and subjective beliefs about physical and economic variables. A strategy consists of many choices (as described above), and these choices taken together describe a farm system.

A dairy farm can be loosely separated into two basic activities (see figure 1). Firstly, the farmer uses some proportion of his or her land to grow forage. Secondly, the farmer combines home-grown forage and purchased feed to offer dairy cows which use it to produce milk. Economic variables such as milk, land, labour, capital and feed prices combined with environmental factors like climate and water availability will determine the optimal arrangement and intensity of the two basic activities. For example, in New Zealand, a high cost of purchased feed and low milk price lead to dairy systems where most land is used to grow grass, which is in turn grazed as the primary feed for the dairy cows. Alternatively, in the US, a relatively low purchased feed price and high milk price lead to farm systems that use larger amounts of capital per cow to ensure higher intake and hence higher production per cow than grazed pasture would allow.

Variability of profit for a dairy farm system is caused by both economic variables and physical interactions. Economic variables include the milk price, the price of purchased feed and appreciation of land values. Physical impacts on pasture growth

---

[1] mneal@student.usyd.edu.au

can be caused by many aspects of climate including low or excessive rainfall, temperature and solar radiation.

There are several aspects of dairy farms in general that make prescriptive analysis of the decision-making process more complex. Firstly, dairy farms are usually family owned operations with the majority of the labour supplied by the owners. The Dexcel Economic Survey (2003) shows 63% of the value of labour is supplied by family members. It appears that many dairy farmers may accept lower wages than an equivalent paid labour unit. While there is an established labour market, many farmers are reluctant to sell their excess labour to other farms, or hire additional labour.

Figure 1: Schematic representation of a dairyfarm

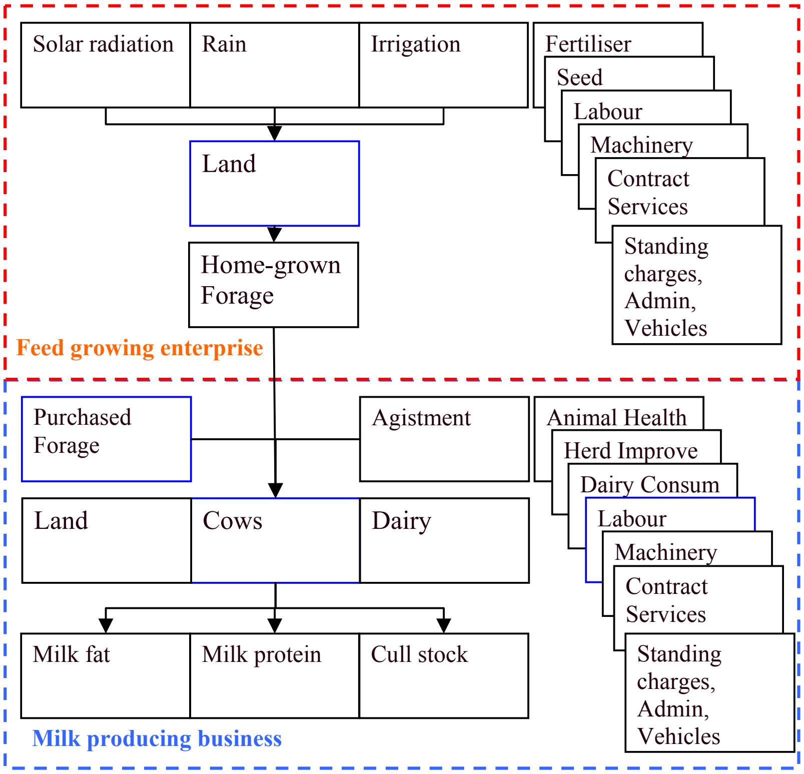

Secondly, the majority of equity has also been provided by the family owners. In 2002-03 this amounted to 61% equity of the average farm, where the average farm was valued at NZD$2.6 million (Dexcel, 2003). It appears that many farmers are also willing to accept less than the equivalent market return for investments of a similar risk profile (Neal, 2004a).

The farmer's decision problem can be simplified somewhat by assuming that a farmer will hire labour if required, or sell labour if there is an excess of labour[2]. Under this assumption the level of labour used by the farm system will not be important in

---

[2] An alternative way of viewing the farmer's provision of labour can allow the same analysis. This would be to assume that although the farmer doesn't hire or sell excess labour, their marginal value for leisure time is equivalent to the market wage, and so they are indifferent between providing any level of labour that the farm requires. However this may lead to logical inconsistencies when considering a utility function implying risk aversion. A second (far simpler) alternative is to assume that the farmer hires all labour.

deciding on a farm system. Hence the farmer's major concern will be the return achieved on their equity. Assuming that a farmer's choice is some convex combinations of a farm system (investment) and a risk free rate (for borrowing or lending) simplifies the problem further[3]. A final assumption that the farmer's utility function can be described in terms of the mean and variance of returns, and this allows the application of separation theorem[4]. The result from separation theorem is that there is only one optimal farm system, regardless of the farmer's level of risk aversion. The farmer responds to their level of risk aversion by choosing a level borrowing or saving to combine with their investment.

Figure 2 illustrates two farmers' preferences expressed in utility function U_A and U_B. The risk free rate of borrowing and saving is shown as R_f, and three possible farm systems (mutually exclusive investments) are shown as I, II and III. Without perfect capital markets Farmer A might prefer farm system I, but under perfect capital markets prefers to save some proportion of equity at R_f, and invest the remaining equity in investment II. Without perfect capital markets Farmer B might prefer farm system III, but under perfect capital markets prefers to borrow some proportion at R_f, and invests the loan proceeds with the equity in investment II.

Figure 2: Two farmers choosing a farm system

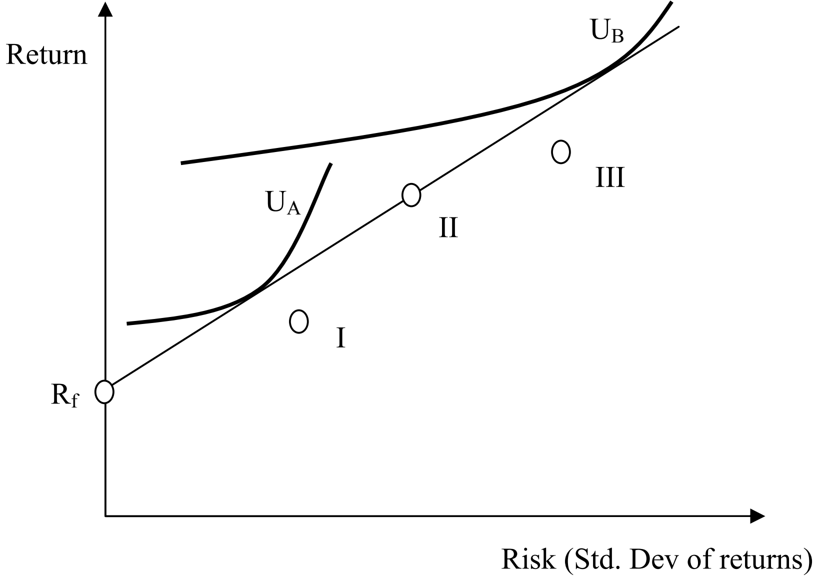

It can be shown that the slope of the line from R_f is the Sharpe ratio. That is, the slope of the line is:

$$Sharpe = \frac{E(R_{II}) - R_f}{\sigma_{R_{II}}} \qquad \text{Equation 1}$$

---

[3] Borrowing and saving at a risk-free rate implies perfect capital markets (no transaction costs). An implicit assumption is constant returns to scale. This is shown to be reasonable for the farms in question (Neal, 2004b). A further assumption is the normal distribution of returns.

[4] A utility function described in terms of mean and variance implies risk aversion but also implies increasing absolute risk aversion (IARA), not a well accepted assumption (Hardaker et al, 2004). However Simmons (2002) suggests that many useful examinations with this assumption can be generalized.

The final result is that the farm system that results in the maximum Sharpe ratio is the farm system that should be chosen by the farmer, regardless of risk aversion characteristics, given the assumptions outlined above.

An alternative way to select a farm system without making as many assumptions is to use stochastic dominance. This method allows a full distribution of outcomes to be compared (removing the restriction that returns be normally distributed). It does not necessarily assume a risk-free rate (although this can be incorporated if required). First degree stochastic dominance makes a fairly mild assumption about the utility function: More is preferred to less.

In mathematical terms, where $F_A(x)$ and $F_B(x)$ are cumulative distribution functions describing the returns from two farm systems A and B, first degree stochastic dominance (B dominates A) can be described by (Hardaker et al, 2004):

$$F_A(x) \leq F_B(x) \text{ for all } x \qquad \text{Equation 2}$$

While first degree stochastic dominance makes few assumptions, it may not be particularly discriminative. In other words, it may not significantly reduce the number of alternative farm systems that the farmer must choose from. This is because a common result from the first degree stochastic dominance test is that neither A nor B dominate the alternative.

Second degree stochastic dominance makes the additional assumption that risk aversion exists. This rule (B dominates A) can be summarised by (Hardaker et al, 2004):

$$\int_{-\infty}^{x^*} F_A(x)dx \leq \int_{-\infty}^{x^*} F_B(x)dx, \text{ for all values of } x^* \qquad \text{Equation 3}$$

This rule reduces the number of alternative farms systems that the farmer must choose between further than for first degree stochastic dominance, although there may still be several farm systems to choose from.

Levy (1992) provides an example of the discriminatory power of the rules described above with reference to 73 mutual funds. He found that using a rule equivalent to the Sharpe ratio reduced the efficient set to one fund. First degree stochastic dominance still had an efficient set of 68 funds (only 7% of funds were dominated). Under second degree stochastic dominance, the efficient set was 16 funds (78% of funds were dominated).

2. Background of Dexcel’s Whole-farm Model

Dexcel is the research and extension arm of New Zealand's dairy industry. The goal of Dexcel is to improve the competitiveness and profitability of New Zealand dairy farmers. The research division within Dexcel has a specialist programme on farm systems. The farm systems programme has the goal of “developing farm systems that optimise resource use for maximum profit and providing farmers with tools and knowledge to manage their resources to best achieve their personal business objectives” (Dexcel, 2004). Within this programme is a strong modelling component and the primary tool used for modelling is the Whole-farm model (WFM).

The WFM is a model of a pasture based dairy farm, implemented using an Object-Oriented (OO) approach using the Smalltalk language. The OO approach allows the incorporation of sub-models that may have been developed elsewhere, including cases where they may be in different programming languages. The WFM has a choice of pasture models; currently a simple seasonal average growth rate model (SimplePasture) and a more complex growth model called McCall, based on the work of McCall et al (2003).The WFM also has a choice of animal models; currently a simple energetics based cow model (SimpleCow) and a more complex model (MollyCow) based on the work of Lee Baldwin (1995) at the University of California, Davis.

Figure 3 shows a simplified schematic of the Whole-farm model. The WFM creates multiple instances of a cow based on the selected animal model and user descriptions. All cows may be different in any physical aspect such as weight, genetic potential or calving date. The WFM creates multiple paddock instances based on the user descriptions. Each paddock may be different in size, although currently the model does not use spatial characteristics of location as an input into the model. The paddocks currently model only the predominant pasture type of Ryegrass. The management policies interact with the cows and paddocks on a daily time step to simulate the biophysical output. This is then used in the economics component to model a simplified profit and loss statement, balance sheet and calculate return on assets.

Figure 3: Simplified schematic of the Whole-farm Model

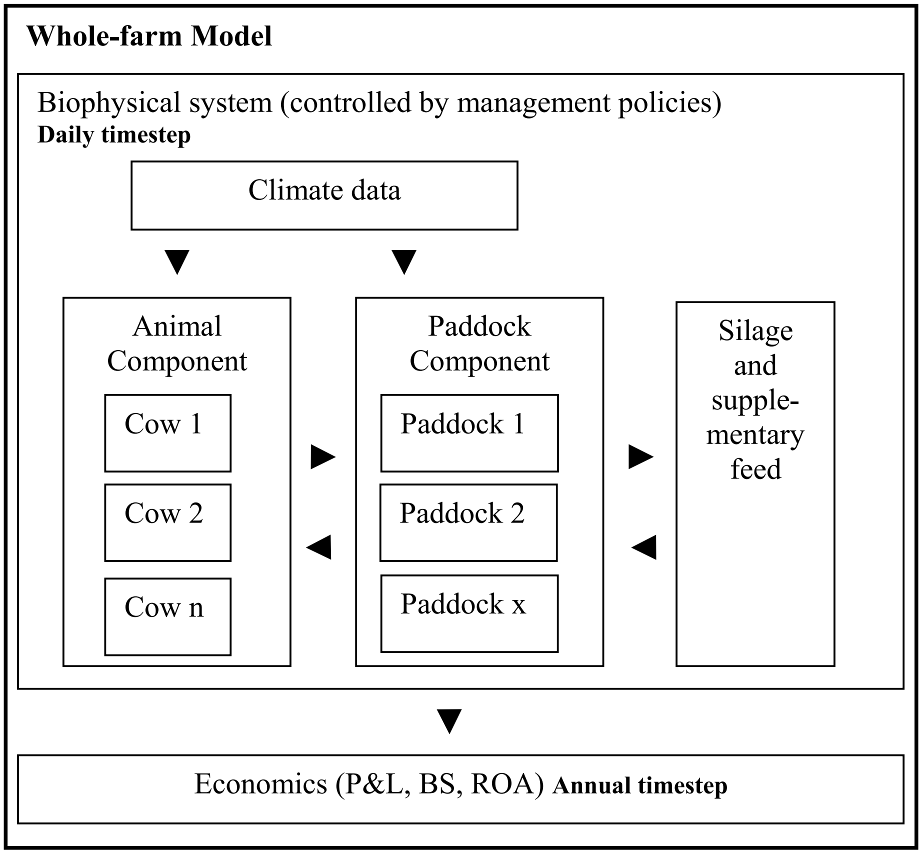

The McCall pasture model describes the dynamics of pasture growth mainly through the following equation where dG/dt is the rate of growth in kg Dry Matter (DM) ha^{-1}.

$$\frac{dG}{dt} = \underbrace{\alpha I.g_t.g_T.g_w.c(G)}_{\text{New growth}} - \underbrace{\sigma_t.S_w.G}_{\text{Senescence}}$$

Equation 5

The first component of new growth consists of incident radiation (I) multiplied by the efficiency of photosynthesis (α). The second component (g_t) modifies growth dependent upon the time of the year, and the third component (g_T) modifies growth dependent upon the temperature. Fourthly g_w modifies growth based on the soil water and finally c(G) simulates the effect of green canopy as a function of green pasture mass (G). The first component of senescence[5] (σ_t) is the base senescence rate and is dependant on the time of the year. The second component (S_w) increases the base senescence with reduced levels of soil water. The McCall model also models the dead components of the pasture as soluble and non soluble. A great deal of the detail of the model can be found in McCall et al (2003).

Molly is a dynamic, mechanistic model and is composed of differential equations describing the metabolism of nutrients. In the WFM these differential equations are integrated over a daily time step. Molly predicts energy metabolism, nitrogen flow, milk production, amino acid metabolism and body composition changes. Molly has been parametised for cows under New Zealand conditions as described in Palliser et al (2001). Feeds are broken down into 18 components (fractions) as described in figure 4.

Figure 4: Fractions in feed used by Molly model

Cellulose	Soluble carbohydrate	Soluble protein
Hemicellulose	Starch	Insoluble protein
Lignin	Organic acids	Non-protein nitrogen
Soluble ash	Pectin	Urea
Insoluble ash	Lactate	
Lipid	Butyrate	
Added fat	Acetate	

The model uses Michealis-Menten type equations to describe the velocity (V_0) at which nutrients react (are metabolised) from an Enzyme (E) and substrate (S), to form a product (P), as shown by the reaction formula (equation 6).

$$E + S \underset{K-1}{\overset{K1}{\rightleftarrows}} ES \xrightarrow{K2} E + P$$

Equation 6

Equation 7 is the Michealis-Menten equation. This equation shows the current velocity (V_0) as a function of the maximum velocity (V_{max}) of the reaction, the substrate concentration [S] and a constant (K_M). K_M is derived in equation 8 from the initial reaction rates.

---

[5] Senescence is the death of plant material.

$$V_0 = \frac{V_{\max}.[S]}{K_M + [S]}$$ Equation 7

$$K_M = \frac{K_{-1} + K_2}{K_1}$$ Equation 8

Overall, the original Molly model defines around 170 state and algebraic parameters and variables to predict the productive behaviour of the cow.

The tactical decision-making of the model is carried out according to management policies. The key policy areas relate to:

1. Pasture feeding;
2. Paddock usage;
3. Supplement feeding;
4. Pasture treatment; and
5. Cow management.

There are multiple policies to choose from, and many policies have default parameters that can be changed by the user. Many policies are related to the decision rules described in detail by MacDonald et al (1998).

The pasture feeding policy is typically set for the cows to meet demand from pasture if available, but the paddock usage policy may limit the proportion of the farm that can be grazed according to typical growth patterns for each season. For example, in winter when pasture growth is typically low, cows are restricted to grazing a proportion of the farm to ensure pasture consumption does not dramatically exceed the growth. In the case where cows would demand more pasture than is fed, this gap may be filled with supplementary feed, depending on the supplement feeding policy.

The pasture treatment policy uses management rules to apply fertiliser and, if available, irrigation. Other management rules apply to cutting grass to make silage, which becomes available to feed out in later time periods (after taking into account harvest and feed-out losses). The cow management policy relates to mating, culling and drying off decisions. Also included under the cow management policy is a choice of milking frequency (once or twice daily) and the ability to agist stock (graze on non-farm areas) at certain times of the year.

The economics component consists of a simplified profit and loss statement, balance sheet and return on assets. The profit and loss statement is developed in a similar format to the Economic Farm Surplus (EFS) (Dexcel, 2004). Revenue is primarily generated through MilkSolids (MS) sales, where this is MS production multiplied by the price per kg MS. Additional revenue is earned from the sale of cull stock, which is assumed to average $180 per cow calved.

A large proportion of costs are defined in an activity based costing framework with default values generated through the use of economic survey data specific to an area. The main cost drivers appear to be the number of cows calved and the effective farm area. Figure 5 shows the costs calculated using the activity based costing. With further improvements in the model allowing management to make decisions that would impact these costs, new cost drivers can be incorporated. Constant returns to scale is assumed and Neal (2004b) found this to be a reasonable assumption over a wide range of farm sizes encompassing more than half of farms surveyed in the 2002 economic survey (Dexcel, 2003).

Figure 5: Cost drivers and associated costs

Cost driver	
Per Cow Calved	Per Effective Hectare
Costs	
Wages	Crop and regrassing
Animal Health	Freight
Herd improvement	Fertiliser
Farm dairy	Weed and Pest
Electricity	Repairs and Maintenance
	Vehicle
	Administration

Feed conserved on farm incurs a cost of $100 per tonne, approximated as the cost of a contractor service. The cost of supplementary feed is calculated based on the reduction in the stock of supplementary feed from the beginning of the year until the end, priced as if additional feed was purchased. A build-up in feed stocks beyond the starting level is valued at the equivalent market price, where this is the purchase price less cost of transport. Adjustments are also made for changes in cow body condition score. Losses in body condition are costed to the equivalent of providing supplementary feed to bring the cow back to the original condition. Gains in body condition score are valued at the feed spared based on the value of purchased feed[6]. Changes in the level of pasture are also adjusted in a similar way. It is possible for a management strategy to lead to the death of a cow and it is assumed to be replaced at the end of the season at the market cost of a new cow. This is a relatively large expense and implies that profitable farm systems do not plan to have cows dying as part of a strategy.

The balance sheet is developed using the broad categories outlined in figure 7, with default appreciation (depreciation) rates. Land typically accounts for half or more of the asset value of the farm, and appreciation rates have averaged over 9% for the past decade, although Neal (2004c) questions the likelihood of this trend continuing given the longer term trend is a real average appreciation rate around 4%.

Figure 7: Asset classes, default values and appreciation rates

Asset class	Default value	Appreciation rate
Land	$18,000 per Ha	4%
Shares	$5.40 per kg MS	10%
Dairy	$3,750 per ha	-6%
Machinery	$1,000 per ha	-20%
Cows	$700 per cow	2%

The WFM attempts to predict what effect variables outside the farmer's control will have on his return, and the risk associated with that return. The major causes of risk are assumed to be:

---

[6] The feed required or spared by a change in condition score takes into account the relative efficiency with which energy is mobilized or deposited

1. Weather – e.g. the risk of a dry or wet year occurring;
2. Milk price – e.g. risk of milk price being higher or lower than the long term average;
3. Supplementary feed price – e.g. risk of the feed price being higher or lower than average; and
4. Capital appreciation rates – e.g. in particular, the risk of land prices moving up or down.

Effectively the risk report is taking a farm system and performing Monte Carlo simulation to find the distribution of returns with variable weather, milk price, supplementary feed price and capital appreciation rates.

The weather risk can only be calculated when a WFM simulation is set up with a component model that relates production to the weather that occurred, i.e. McCall pasture model. Any year for which climate data is available can be used in the Monte Carlo simulation, and simulated weather could also be used. Because weather impacts on the amount of pasture produced in a whole region, it would also be expected to impact on the price of supplementary feed. This effect is modelled by varying the mean supplementary feed price with the potential pasture production of the climate year.

The price risk (ie milk price, supplementary feed price and asset prices) is found using the distribution of prices determined from actual data. These prices are assumed to be normally distributed (Neal, 2004d). Sets of prices are generated from uniform random numbers and transformed into normal distribution by means of an inversion process (*ppnd* algorithm as presented by Beasley and Springer, 1977) to give a matrix Z. Z in this case is a matrix of *n* price sets (default of 100) by *p* prices (default of 3)

The random (uncorrelated) prices are then correlated using the process described by Iman and Conover (1982). This process involves a user defined correlation matrix, C which is a symmetric matrix of size *p*, where *p* is the number of variables (prices) to correlate. To be valid, the correlation matrix must be positive definite. This can be checked by ensuring the smallest eigenvalue of C is positive. If this is not the case, C can be adjusted to become positive semidefinite and as close as possible to the user defined correlation matrix. The smallest eigenvalue of C is found and labelled E_0. Then an adjusted C matrix, C' is found:

$$C' = C - E_0.I \qquad \text{Equation 9}$$

The new (valid) C matrix, C'' is then:

$$C'' = \frac{1}{1-E_0}.C' \qquad \text{Equation 10}$$

Hart et al (2003) suggest using a rank correlation rather than Pearson correlation as this allows a distribution free approach of imposing correlation. Distribution free implies that the distribution of each variable need not belong to a particular family (eg normal). Z is the uncorrelated[7] matrix of prices found above. R is a matrix of (van der Waerden) scores, ranked in the same way as the elements of Z[8]. C'' is the target (symmetric) matrix of (rank) correlations, but D is the current rank correlation matrix

---

[7] Not intentionally correlated. A "random" or non-zero correlation may exist. The variance reducing method is used here to ensure the final matrix of prices is as close as possible to the desired correlation.

[8] N Van der Waerden scores are created using the formula $R_{ij}=\mathrm{Erf}(i/(N+1))$

of Matrix R. V must be calculated as the Cholesky decomposition of C'' and U as the Cholesky decomposition of D[9]. The following matrix multiplication is the performed:

$S=VU^{-1}$ Equation 11

$R^{*}=RS$ Equation 12

Sorting X to give the same ranking as R* gives X*, a matrix correlated according to the desired matrix C''. This matrix of prices can then be used in Monte Carlo simulations using the WFM.

3. Optimisation with an Evolutionary algorithm

Optimisation of the WFM was performed with a specific evolutionary algorithm – a variant of more common genetic algorithms. Heitkotter and Beasley (2004) explain that an "evolutionary algorithm is an umbrella term used to describe computer based problem solving systems which use computational models of some of the known mechanisms of evolution as key elements in their design and implementation." There are a variety of areas of research that fall under this area, including:

1. Genetic algorithms (GA);
2. Genetic programming ;
3. Evolutionary programming;
4. Evolution strategies; and
5. Classifier systems.

The key aspects of similarity between these areas is that optimisation is through the use of a population of possible solutions, and aspects of biological evolution, such as reproduction, selection, mutation and recombination are used to improve the solutions.

Genetic algorithms have been used extensively across many disciplines including economics and engineering and used for many purposes including scheduling and design problems. Genetic algorithms are particularly suited to complex problems where non-linearity exists and stochastic error occurs in the measurement of results. However, unlike closed form analytical solutions, it is generally not possible to show that a genetic algorithm has produced the optimal solution.

In simple terms, the problem can be defined as selecting the inputs for a function (X) to provide the 'best' solution (Y), where the function is described by:

$Y=f(X)+\varepsilon$ Equation 13

The function may have multiple inputs and multiple outputs, where X would be a vector of n inputs, and Y could be a vector of m outputs. For the decisionmaker, preferences may be determined by the utility (U) derived by the vector of outcomes (Y). Specifically:

$U=g(Y)$ Equation 14

---

[9] ie $D=UU^{T}$ and $C''=VV^{T}$

Thus the overall problem can be simplified to the maximisation (or minimisation) of utility, where utility describes the fitness of the inputs. In other words, fitness is the value of utility, given the inputs (X).

$$U=h(X) \qquad \text{Equation 15}$$

However, due to the original function (f) that includes non-linearity, discontinuity and/or stochastic measurement error, no closed form analytical solution exists. This maximisation problem is represented with individual solutions (individuals) in figure 8.

Figure 8: Representation of the maximisation problem

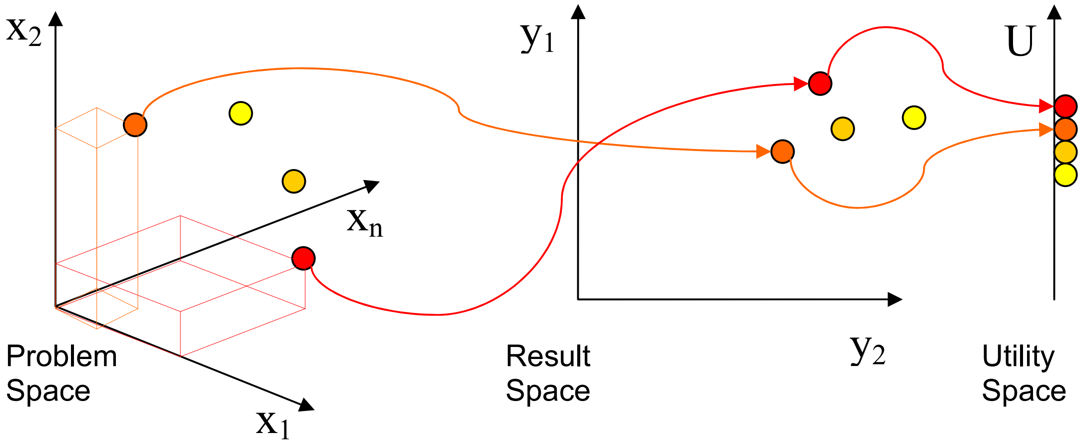

The initialisation of a genetic algorithm begins with the generation of a population of solutions. The population consists of a number of individuals, and each individual is described by a chromosome. The chromosome encodes all the information given by the values of the inputs ($X=\{x_1, x_2,\ldots,x_n\}$). An allele is the value at a certain point (locus) of the chromosome. Figure 9 shows a graphical representation of a population as three binary coded individuals.

Figure 9: A population of individuals

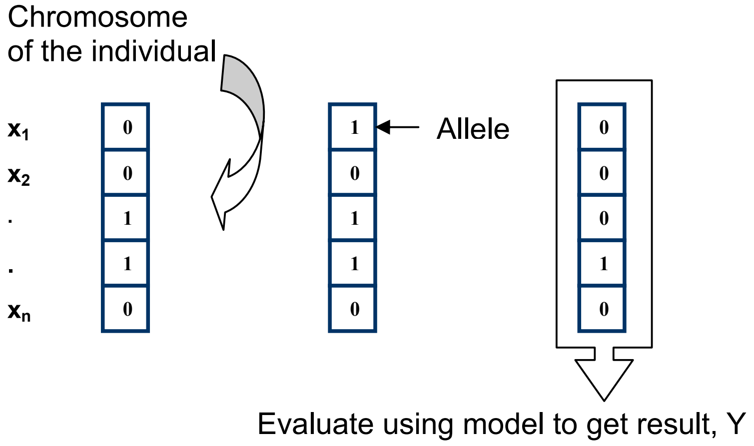

The encoding of the chromosome may be using real numbers or as binary values, although real values have become more popular, partly for efficiency reasons as outlined in Liu (2003). The number of individuals to have in a population has been an aspect of debate and the most effective value will depend on the problem and the type of genetic algorithm. However Mayer et al (2005) suggests between 1.5 to 2 individuals for each input variable (x_i) may be adequate. Clearly a larger population creates more diversity, but uses more resources for computation.

After the initial population has been generated, each individual must be evaluated by the function (h(X)) to give the fitness of the individual (U) - the measure of how well the individual meets the objective. Evaluation of the original function (f(X)) is typically the most computationally intensive part of the process. The evolutionary operators of selection, recombination and mutation are then typically used. The selection operator selects two parents from the original population for mating, based in part on their fitness. In other words, better solutions have a higher probability of passing on their genes to the next generation. The most popular selection operators are roulette or tournament selection.

Roulette selection is similar to spinning a roulette wheel where each individual has a probability of selection proportional to their fitness. For example, ‘high’ fitness individuals might be twice as likely as ‘low’ fitness individuals. This system can also work with fitness ranking rather than the actual fitness level. Liu (2003) uses this approach and the probability of selection for the individual ranked *q* is:

$$\mathrm{P}_q = \frac{(N-(r_q-1))^s}{\sum_{i=1}^{N}(N-(r_i-1))^s} \qquad \text{Equation 16}$$

The parameter *s* is the selection pressure, and increasing values of this improves the probability of selecting individuals ranked highly.

Tournament selection involves the selection of two or more individuals who ‘compete’ based on their fitness value. The ‘winner’ is used for mating (recombination). Clearly the worst individual in the population will never be used for breeding (Drake et al, 1998).

Usually two parents are selected to breed a single new individual, although there are variations on this theme. The main mating (recombination) operator is crossover. A point in the chromosome is randomly chosen to cut the chromosomes of the parents. The child is then formed by concatenating a chromosome head with a tail, as illustrated in figure 10. Double point crossover is only slightly more complex (see figure 10), although there are many other crossover implementations.

Figure 10: Mating two parents; single point crossover, double point crossover

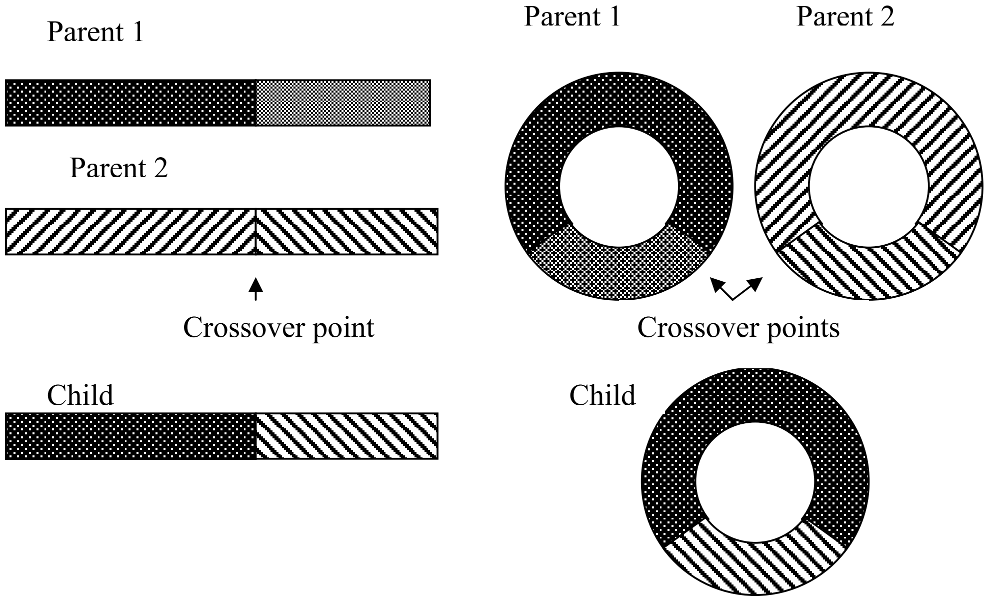

After a child has been successfully bred, a mutation operator is generally applied. This is generally introduced with a low probability to introduce new alleles into the population. The implementation by Liu (2003) uses a mutation probability (default of 0.15) to decide whether an allele will be altered. If a mutation occurs, a random amount is added to the value of the original allele. The random amount is normally distributed with a mean of zero and standard deviation equal to one tenth the width of the zone[10]. Mutation can also be introduced by means of a population gap. This is where the poorest performing proportion of the population are discarded and replaced by randomly created individuals.

When new individuals have been bred they are evaluated in order to create a new population. Elitism is the feature describing how a new individual becomes part of the new population. Some GA's implement total elitism, where parents and children are grouped and only the individuals with the highest fitness become the new population. Alternatively, a GA may implement partial elitism where a child only has to have higher fitness than one of the parents to become part of the new population. Total elitism has a positive impact on convergence, but reduces diversity. The process of the genetic algorithm is represented in figure 11.

---

[10] The width of the zone is the maximum value of the allele x_i minus the minimum value of x_i

Figure 11: The genetic algorithm process

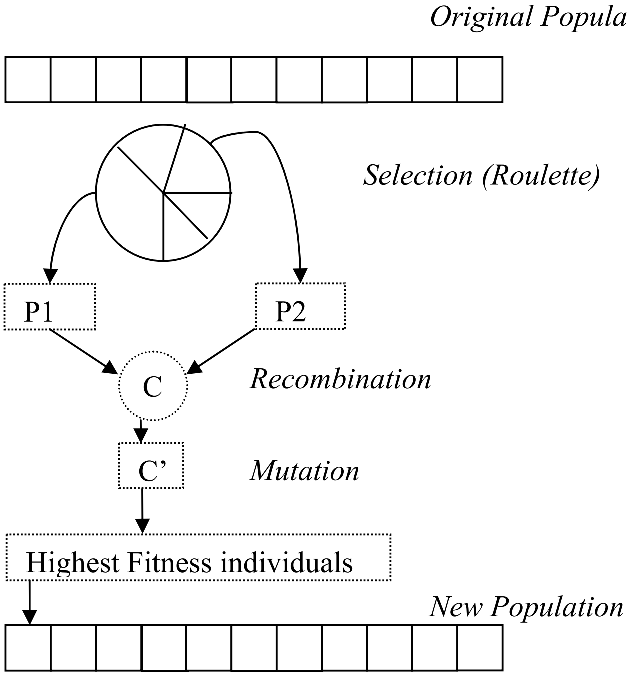

The acceptance of a new generation will then result in a further process of selection, recombination and mutation. This loop continues until a termination criterion is satisfied. Termination criteria could include a maximum number of generations; some predetermined level of convergence in the chromosomes; or some minimum progress in the fitness value.

The evolutionary algorithm used for the optimisation of the Dexcel WFM is a variant of the typical genetic algorithm just described. This differential evolution algorithm is implemented in a similar way to Mayer et al (2005) based on the work of Storn and Price (1997). The important features are that it performs recombination and mutation as a single step, and the way this is implemented leads to the benefits of adaptive search. Adaptive search implies that knowledge of the diversity of the current population is used in creating new individuals.

Firstly a member of the population is chosen as a Parent (P). A child (C) is then created to compare to the parent, and replaces the parent if it has a higher fitness level. To create the child three members are selected at random from the population[11]. These can be labelled X, Y and Z. The difference between the allele values of X and Y produce a vector of values (D) that represent a random measure of the diversity in the population. A proportion of this difference is then calculated using a scaling factor (f) multiplied by the vector of differences (D) to create (E). The Vector E is then added to Z to create another parent (G). The child is then created by choosing an allele from the parent P with probability CR, or the respective allele from the parent with probability (1-CR), where CR is the crossover probability. This process is shown graphically in figure 12. Mayer et al (2005) also implements a useful method suggested by Kinghorn (2005) to prevent premature convergence due to the

---

[11] Selection could be done with roulette or tournament selection, but this implementation is random.

interpolative nature of using an f value of less than one. The suggestion is to every n generations “pulse’ f to a value (much) larger than one to facilitate extrapolative search.

Figure 12: Representation of the differential evolution process

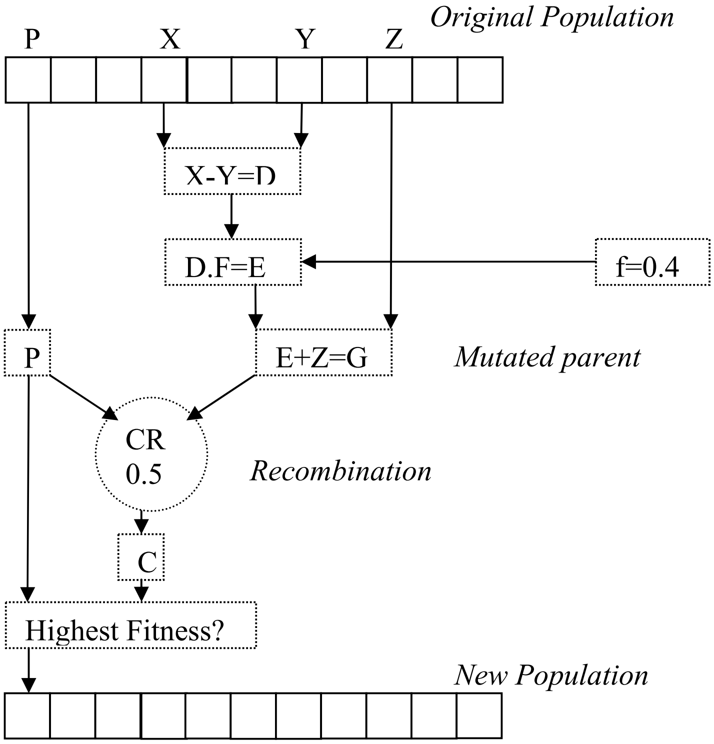

Adapted from Storn (2004)

In the WFM, selection is driven by fitness (U=h(X)), where this can be implemented in one of two ways. Firstly, the Sharpe ratio was used. This rule is very discriminating, as explained in Section 2. Secondly, stochastic dominance (first order or second order) was used. A common outcome of a stochastic dominance comparison is that no dominance relation can be found. In this event, the algorithm keeps the parent for the new population with probability (*Q*), or discards the parent and keeps the child with probability (*1-Q*). If *Q* is set equal to 1, this ensures that no acceptable individual is lost from the population. However a smaller value of *Q* enhances the search capability of the algorithm.

One of the benefits of a genetic algorithm is the implicit ease of parallelising the process. This is because the calculation of the individuals fitness, Y=f(X), which is typically computationally expensive, can be calculated independently of the selection, recombination and mutation process of the GA. This feature is used to maximum advantage in the WFM optimisation implementation as the fitness calculation is distributed to computers in the available network. The Dexcel network consists of around 70 computers that are idle around 70% of the week. The specific implementation is a master-slave approach, where one computer is a master, and runs the GA. The master then passes individual chromosomes to other computers on the network (slaves) to evaluate the fitness. The master-slave arrangement is shown in figure 13.

Figure 13: Master-slave arrangement

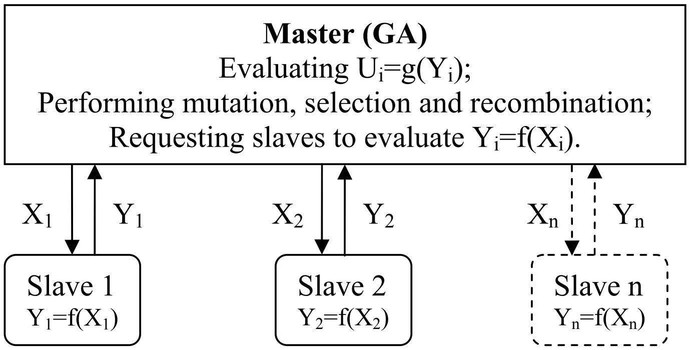

Through the use of secure virtual private networks and remote viewing software, the optimisation can be run from any computer with the correct software installed. Any computer or network of computers can be used as a slave once the software is installed. This is because the communication between slave and master is via generic internet protocols.

4. Results

To generate the results presented here, two optimisations were initiated; the first using fitness comparison based on the Sharpe ratio and the second using first degree stochastic dominance criteria. Both optimisations used the following configuration to find the optimal farm system(s)[12]. Four inputs were used to describe the farm system:

1. Calving Date, bounded between 2 June and 31 August
2. Dry Off Date, bounded between 16 March and 14 June;
3. Stocking Rate, bounded between 1.5 and 6 cows per hectare; and
4. Initial Amount of Silage, bounded between zero and 4.5 wet tonnes per cow;

Each fitness level involved the simulation of two climate years, 1994 and 1995, where the former year showed above average pasture growth and the latter showed below average growth. The population size was set at 20 and the termination criterion was 30 generations. The crossover rate (CR) was set to 0.5. The interpolative factor (f) was set to $f_i=0.4$ and pulsed to $f_e=4.0$ every 4 generations for extrapolative search. The Sharpe Ratio was calculated based on a risk free rate (r_f) of 5%. When the stochastic dominance rule did not find a dominance relationship, the probability for selecting the parent (Q) was 0.5. Each of the two optimisations was repeated three times to check for consistency of results, although the majority of results discussed relate to the first replicate of each optimisation.

Using the Sharpe ratio as the fitness function leads to a larger improvement in the best Sharpe ratio than using stochastic dominance, due mainly to its improved discriminatory power (figure 14). In the stochastic dominance optimisation, around 50% of the replacements were by dominance relationships, giving some idea of its discriminatory power. Interestingly, the average Sharpe ratio of the children of each generation does not appear radically different depending on the fitness function (figure 14). This would be due to the nature of the original function ($Y=f(X)$) and could not be expected to occur in all related problems.

---

[12] With any stochastic optimisation it is not possible to prove that a given result is optimal without evaluating all possible options. Only near optimality is suggested, although this is generally near enough for the purpose. The total number of farm systems spanned in this problem space is 10^8.

Figure 14: Sharpe ratio across an optimisation with two fitness measures.

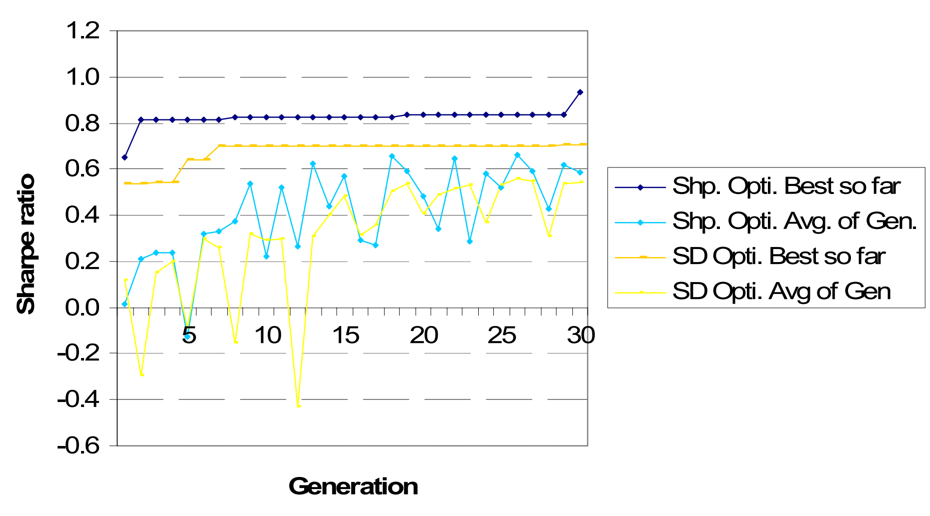

All individuals from both optimisations can be examined in the mean-standard deviation space (figure 15). The Security Market Line (SML) is shown as an upward sloping line from the risk free rate of 5%. The highest Sharpe ratio was found by the Sharp ratio optimisation, but in this replicate only during the last generation. This suggests that more generations of this optimisation would have found even higher Sharpe ratio individuals. The final generation of the stochastic dominance driven optimisation have the appearance of a frontier, but it is well below the highest Sharpe ratio found.

Figure 15: Individuals from optimisations in a mean-standard deviation space

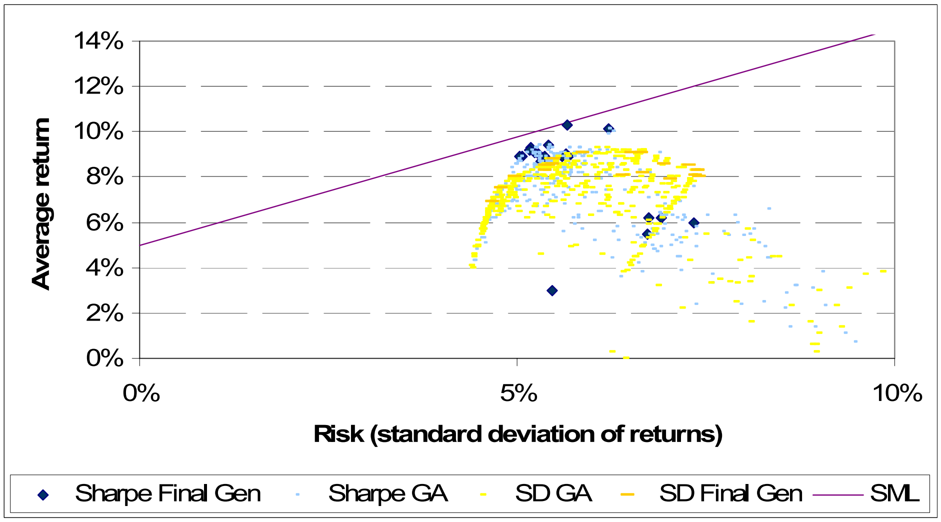

Graphing the cumulative distribution functions of a family of successive individuals provides an insight as to why first degree stochastic dominance finds it difficult to improve over successive generations. Figure 16 shows a tendency towards individuals that are clustered closely and crossover in at least one point.

Figure 16: A family of successive individuals from a stochastic dominance optimisation, represented as cumulative distribution functions.

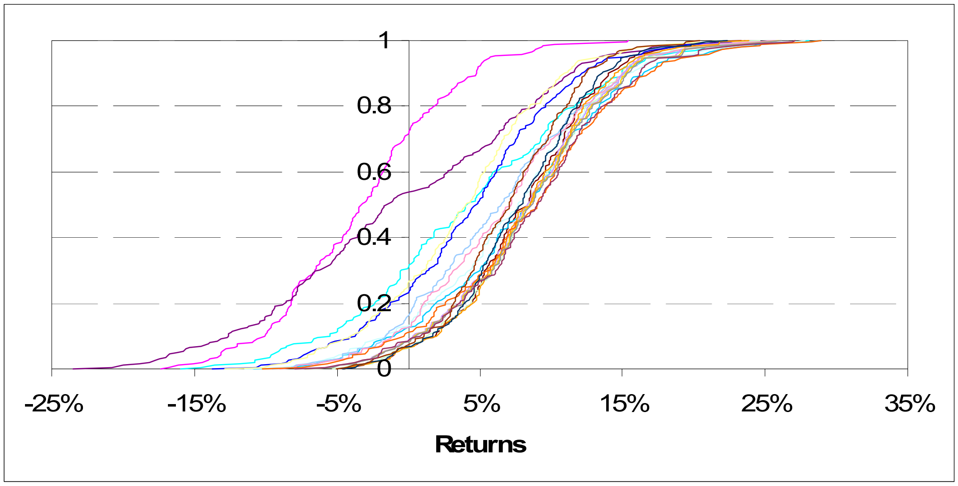

One of the key assumptions allowing the use of the Sharpe ratio is that the returns of the farm systems are normally distributed. The Shapiro-Wilks normality test was tested on each individual evaluated over both the Sharpe and stochastic dominance based optimisations. The null hypothesis is that the returns are normally distributed. Figure 17 shows the p-value for a replicate from a Sharpe and stochastic dominance based optimisation. While there are some individuals where normality can be rejected at the 5% level of significance, the final generation for each optimisation does not contain these individuals. However these results would be sensitive to the assumed distributions of economic variables and possibly the climate years chosen for the simulations.

Figure 17: The P-value from a Shapiro-Wilks normality test on returns.

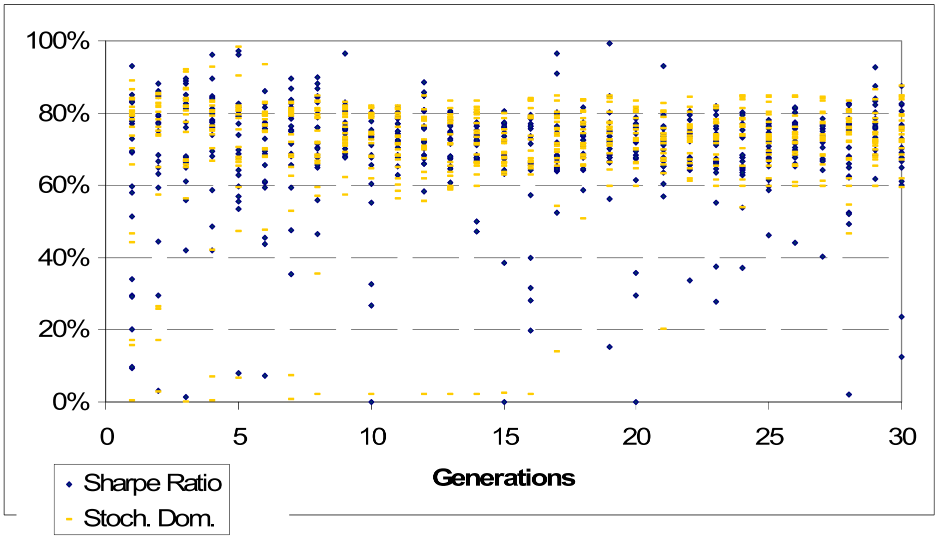

To analyse the results of an optimisation is sometimes difficult due to the multi-dimensional aspect of both the inputs and results. One method to ease graphical interpretation is the parallel coordinate visualisation technique (Inselberg, 1981). This technique represents each variable as a vertical line and joins the relevant values with a line (ie representing an individual of interest). Liu (2003) and Post (2004) created and updated a program (OptVis) to allow the visualisation of any number of individuals for any number of inputs and results in a colour mapped visualisation.

The last generation from the Sharpe ratio optimisation showed most individuals with a moderate stocking rate, late dry-off date, a range of initial silage early calving dates. However as previously mentioned there was large progress made by an individual with the highest Sharpe ratio. This individual was characterised by a very high stocking rate, very early dry-off date and a moderate calving date, and represents a very different type of farm system. The final generation is displayed in figure 18.

Figure 18: Last generation from Sharpe Ratio optimisation

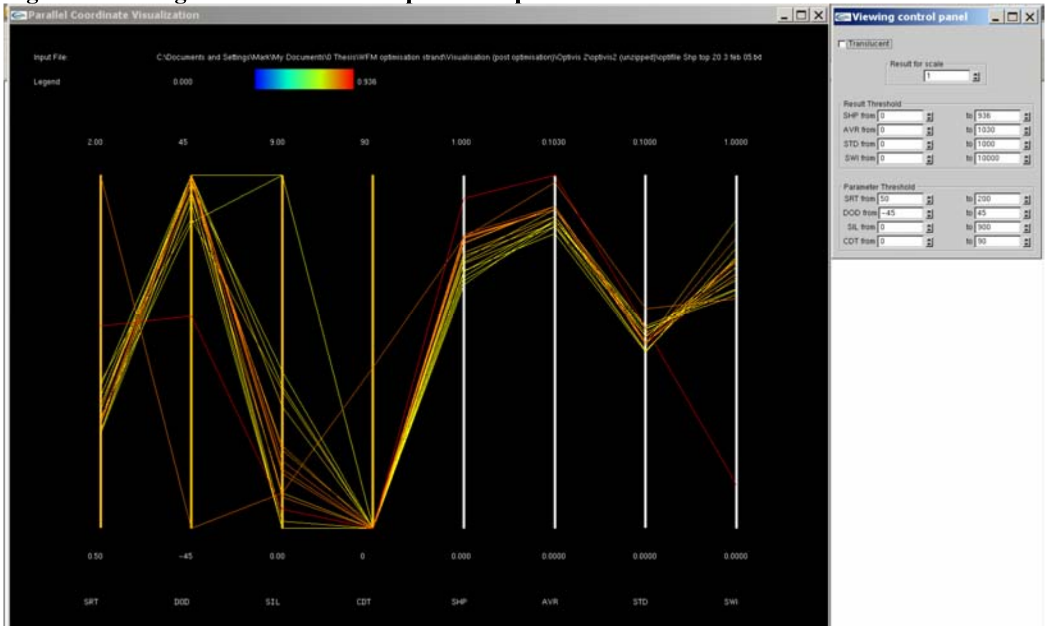

Input File C:\Documents and Settings\Mark\My Documents\0 Thesis\WFM optimisation strand\Visualisation (post optimisation)\Optvis 2\optvis2 (unzipped)\optfile Shp top 20 3 feb 05.txt

Legend 0.000 0.936

☐ Translucent

Result for scale 1

Result Threshold
SHP from 0 to 936
AVR from 0 to 1030
STD from 0 to 1000
SWI from 0 to 10000

Parameter Threshold
SRT from 50 to 200
DOD from -45 to 45
SIL from 0 to 900
CDT from 0 to 90

Four inputs

1. SRT: Stocking Rate, bounded between 1.5 (0.50) and 6 (2.00) cows per hectare
2. DOD: Dry Off Date, bounded between 16 March (-45) and 14 June (45)
3. SIL: Initial Amount of Silage, bounded between zero (0.0) and 4.5 (9.0) wet tonnes per cow
4. CDT: Calving Date, bounded between 2 June (0) and 31 August (90)

Four results

1. Sharpe ratio (0.0 to 1.0)
2. Average return (0 to 10.3%)
3. Standard deviation of returns (0 to 10%)
4. Shapiro Wilks p-value. (0.00 to 1.00)

The final generation of the stochastic dominance based optimisation provided solutions that were quite different to those from the Sharpe optimisation. There were similar tendencies to late dates and early calving dates, but these values were more diverse. A much wider range of stocking rates was still considered, but initial silage levels were quite high. Initial silage levels were probably high due to the insurance effect of having feed on hand during feed shortages.

Figure 19: Last generation from Stochastic Dominance optimisation

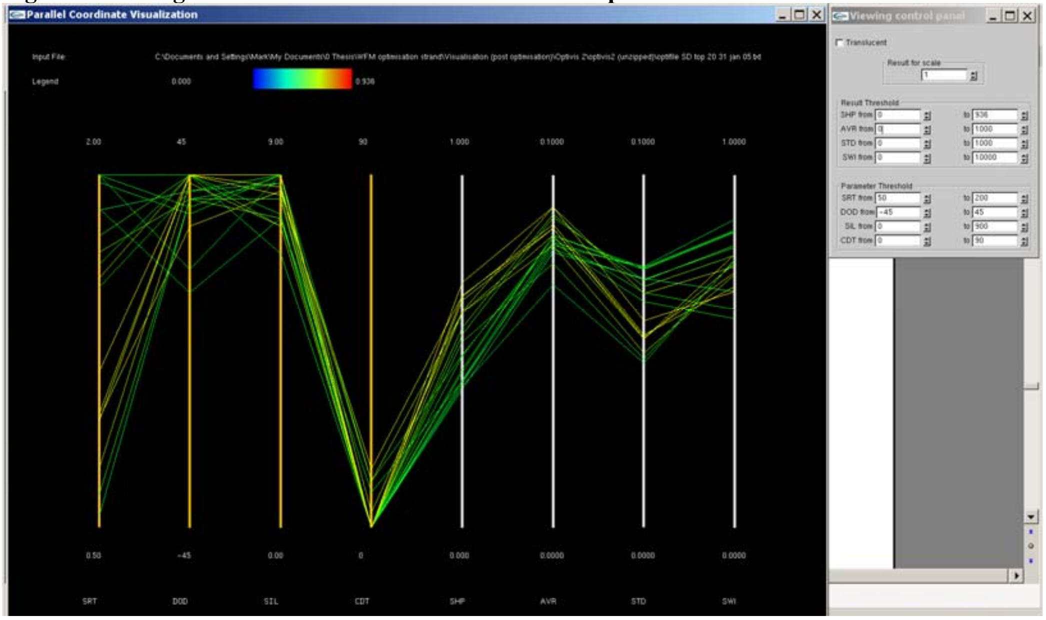

Input File: C:\Documents and Settings\Mark\My Documents\0 Thesis\WFM optimisation strand\Visualisation (post optimisation)\Optivis 2\optivis2 (unzipped)\optfile SD top 20 31 jan 05.txt

☐ Translucent

Result for scale: 1

Result Threshold
SHP from 0 to 936
AVR from 0 to 1000
STD from 0 to 1000
SWI from 0 to 10000

Parameter Threshold
SRT from 50 to 200
DOD from -45 to 45
SIL from 0 to 900
CDT from 0 to 90

Stochastic optimisations by their nature will not produce the same results each time they are repeated. For comparative purposes, 3 replicates were performed and figure 20 shows the progress of the Sharpe ratio for the Sharpe ratio based optimisations. The first replicate was for the results presented above and actually had the lowest Sharpe ratio of the replicates, although only slightly lower than the second replicate. The best replicate was the third replicate and this suggests that the optimisation configuration did not allow for enough searching to ensure near-optimal regions could be reliably found. Solutions to this could include:

i. increasing the number of generations;
ii. increasing the pulse size and/or frequency of extrapolative search (f_e);
iii. increasing the population size.

Figure 20: Comparing 3 replicate optimisations of Sharpe ratio

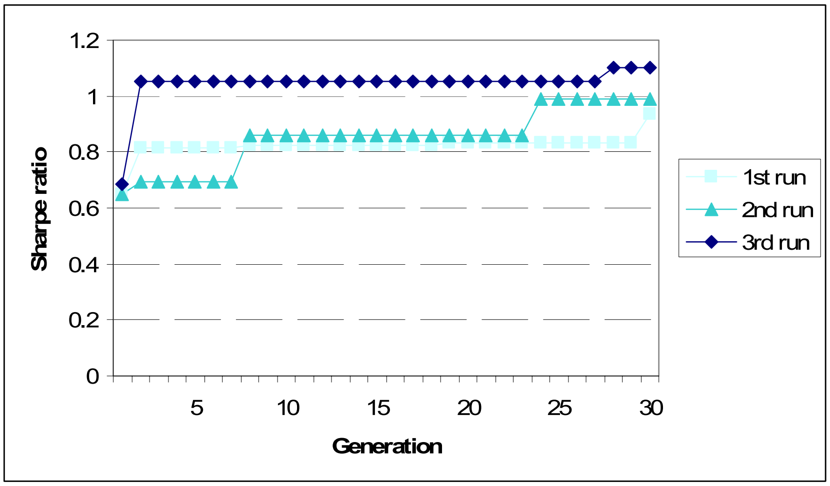

5. Discussion and Conclusions

As a selection tool, the Sharpe ratio proved quite discriminatory. However it may not be applicable to all farmers due to the restrictive assumptions. The reduced assumptions of stochastic dominance may be less restrictive, but the optimisation provides several possible farm systems. The farmer would then have to apply some other criteria to in order to select a farm system.

A major limitation of the results presented is that the replicates showed large differences in the final generations. This was due to the optimisation process not consistently finding “good” regions in the problem space with the current configuration. The results section provided some possible solution to this problem. However the stochastic nature of the optimisation makes it effectively impossible to repeat and compare solutions.

Other basic limitations include the assumed distribution of prices, and the small set of possible climates used in the optimisation process. The model may be a source of error due to factors that are not explicitly modelled such as pugging effects in the McCall pasture model, or aspects not perfectly modelled such as the dry matter intake model of MollyCow. The scaling of the model to an average farm may also cause a minor bias. The economics model did not expressly model the capital requirements for feeding high levels supplementary feed and this would expect to positively bias towards high stocking rates and high use of supplementary feed.

In both the Sharpe and stochastic dominance based optimisations, utility was assumed in some way to be a function of income. Where the ability to borrow and save with well functioning capital markets does not exist, this assumption becomes less tenable. Brennan (2001) provides an example of utility as a function of consumption for developing country farmers.

It is possible that the farmer expresses preferences over outcomes quite different to average income and income variability. For example, a farmer might have large aversion to bankruptcy. The probability of bankruptcy would be related to the liquidity of a farm system, and this depends on cashflows rather than profits. The impact of this might be a bias towards low risk farm systems with reliable cashflows and lower non-cash profits such as land appreciation.

The farm systems are assumed to be mutually exclusive. However it would be possible to own multiple farms, and each farm may not run the same system. To the extent that different farm systems are not perfectly correlated, it would be possible to gain some benefits from diversification. A more likely scenario is farmers running a successful farm system, implementing the same farm system in different areas, although diversification would not be the primary motivation.

Where alternative investments are available it would be possible to create a portfolio including some farm investment and share market investments. To the extent that they are less than perfectly correlated, some diversification benefit exits. However for most farmers, the farm comprises more than 90% of the investment portfolio. The low use of alternative investments is due to a number of factors such as familiarity, perceptions of risk and perceived economies of scale.

When a farmer is making a strategic decision, other factors typically influence the decision. The labour market in a particular region and the farmer's experience with hiring labour will impact on the choice of farm system. If the farmer currently operates a farm, the size and characteristics of the land, and the current animal genetics and calving pattern will affect the choice of strategy. The benefits of moving to a better farm system would be reduced by the cost of adjustment.

The current optimisation did not take into account temporal correlation. Temporal correlation can occur through milk price cycles. Although first order autocorrelation was not found by Neal (2004d), it is possible that a conditional heteroskedascity exists. There is some correlation in land prices over time and correlation with milk price cycles. The impact of correlation could be reduced by modelling longer time periods using a historical correlation matrix.

6. Future work

Improvements to the WFM are incorporated regularly and the most recent version is used for optimisation. The short term direction of the current research is towards improved optimisation configuration, directly modelling the capital required for different feeding levels, and the addition of second degree stochastic dominance as a fitness function.

Longer term research could involve multi-objective optimisation to include preferences over labour use, liquidity and environmental outcomes. Further exploration of temporal correlation of prices may also provide insights into farm systems that are less risky over time.

Tactical optimisation could improve farmers' response to the physical and economic environment. For example, should the farmer cull earlier than planned in response to high feed prices, or should the farmer dry off the herd earlier due to long range weather forecasts.

References

Beasley, J.D. and Springer, S.G. (1977) Algorithm AS111: percentage points of the normal distribution, *Applied Statistics*, 26, 118-121.

Brealey, R. A. and Myers, S. C. (2000), *Principles of Corporate Finance,* McGraw Hill, Sydney.

Deb , K. (2001) *Multi-objective optimisation using evolutionary algorithms*, John Wiley and Sons, Chichester.

Dexcel Ltd (2003a), Dexcel – Research – Farm systems, [Online], Available from: URL http://www.dexcel.co.nz/r_farm_sys.html, [Accessed 15 August 2003]

________ 2003b, *Economic survey of New Zealand Dairy Farmers 2001-02*, Dexcel, Hamilton.

Drake, A. E. and Marks, R. E. (1998) *Genetic algorithms in economics and finance: Forecasting stockmarket prices and foreign exchange* – A review, Australian Graduate School of Management Working Paper, 98-004, 1-32.

Fishman, G. S. (1995) Monte Carlo: Concepts, algorithms and applications, Springer, New York.

Hardaker, J. B., Huirne, R. B. M. and Anderson, J. R. (1997), *Coping with risk in agriculture,* CAB International, Wallingford.

Hart, C. E., Hayes, D. J. and Babcock, B. A. (2003), Insuring eggs in baskets, Center for Agricultural and Resource Development Working Paper 03-WP 339, 1-27

Iman, R. L. and Conover, W. J. (1982) A distribution-free approach to inducing rank correlation among input variables, *Communications in Statistics: Simulation and Computation*, 11, 314-34

Liu, D. (2003), *Techniques for optimisation of a dairy farm simulation model,* Masters Thesis, Lincoln University, Christchurch.

Mayer, D.G., Kinghorn, B.P. and Archer A. A. (2005) Differential evolution – An easy and efficient evolutionary algorithm for model optimisation, *Agricultural Systems*, 83, 315-328

Neal, M (2004a) *What direction to higher returns?* Australian Dairyfarmer, Accepted for publication.

Neal, M (2004b) *Re-ranking of individual firms using alternative performance measures: The case of New Zealand dairy farms*, Proceedings of the Asia-Pacific Productivity Conference, Brisbane, 14-16 July

Neal, M (2004c) *Risk and return in New Zealand Dairy farming*, Dexcel-Link (bulletin series), Autumn 2004, Dexcel, Hamilton

Neal, M (2004d) *Optimising return in a whole-farm model under conditions of risk*, Proceedings of the AARES Conference, Melbourne, 14-16 February

Palisade (2004) *@Risk Manual*, Newfield, New York

Palliser CC, Bright, KP, Macdonald KA, Penno JW, Wastney ME. Adapting the MOLLY cow model to fit production data from New Zealand animals. Proceedings of the New Zealand Society of Animal Production, 61:234-236. 2001.

Post, E. P. (2004) *User manual for OptVis version 2.0*, Christchurch.

Robert, C. P. and Casella, G. (1999) *Monte Carlo statistical methods*, Springer, New York.

Royston, P. (1982) Algoritm AS 181: The W test for normality, *Applied Statistics*, 31, 176-180

Shapiro, S. S. and Wilk, M. B., (1965) An analysis of variance test for normality (complete samples, *Biometrika*, 52 591-611

Sherlock, R.A. and Bright, K.P. (1999) An object-oriented framework for farm system simulation. In Oxley, L. Scrimgeour, F. and Jakeman, A. (eds) *MODSIM99 – Proceedings of the International Conference on Modelling and Simulation*, Modelling and Simulation Society of Australia and New Zealand Inc., Hamilton, NZ. pp 783-788.

Simmons, P. (2002) Why do farmers have so little interest in futures markets?, *Agricultural economics*, 27, p1-6

Spall, J. C. (2003) *Introduction to stochastic search and optimisation*, John Wiley and Sons, New Jersey.

Thode, H. C. (2002) *Testing for normality*, Marcel Dekker, New York.

Wastney M.E., Palliser C.C., Lile J.A., Macdonald K.A., Penno J.W. and Bright K.P. (2002) A Whole Farm Model Applied to a Dairy System. *Proceedings of New Zealand Society of Animal Production* 62:120-123.

Appendix A: Graphs of optimisation progress

◆ Sharpe Ratio
- Stoch. Dom.

Figure A.1

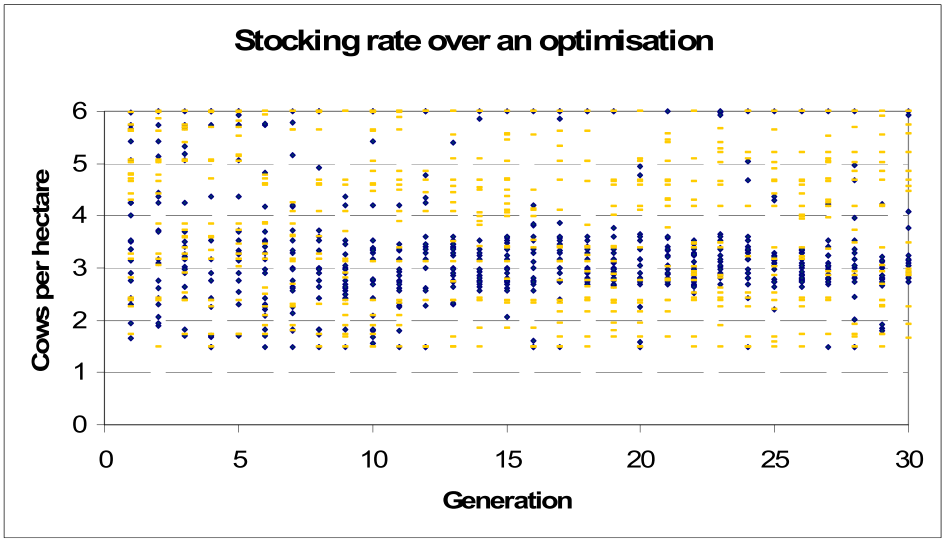

Figure A.2:

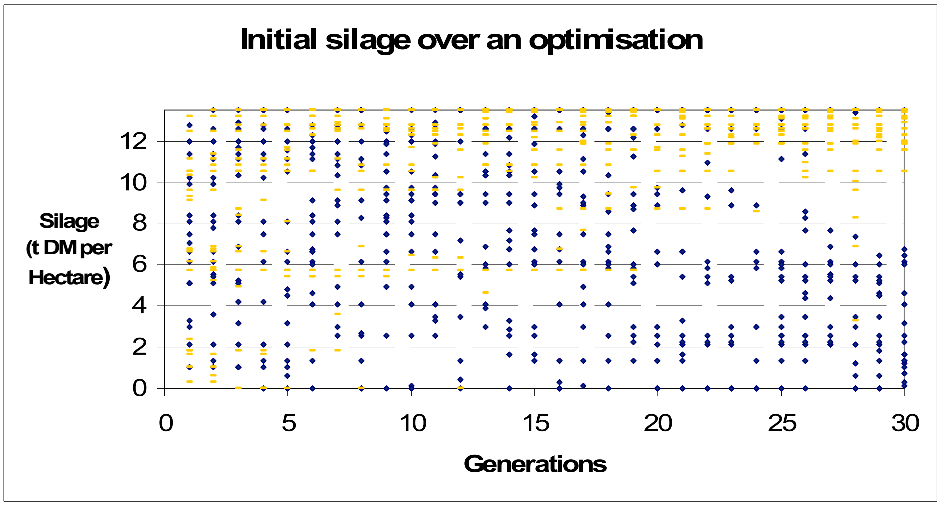

Figure A.3:

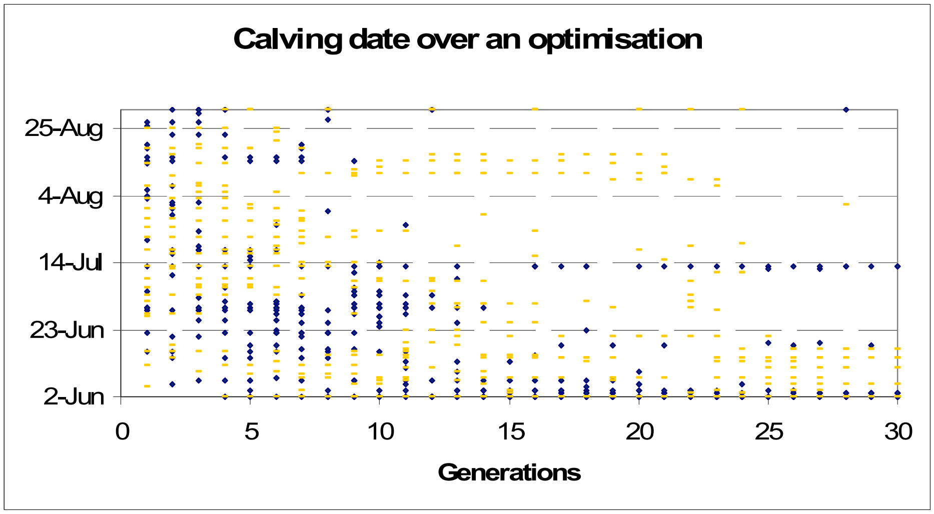

Figure A.4:

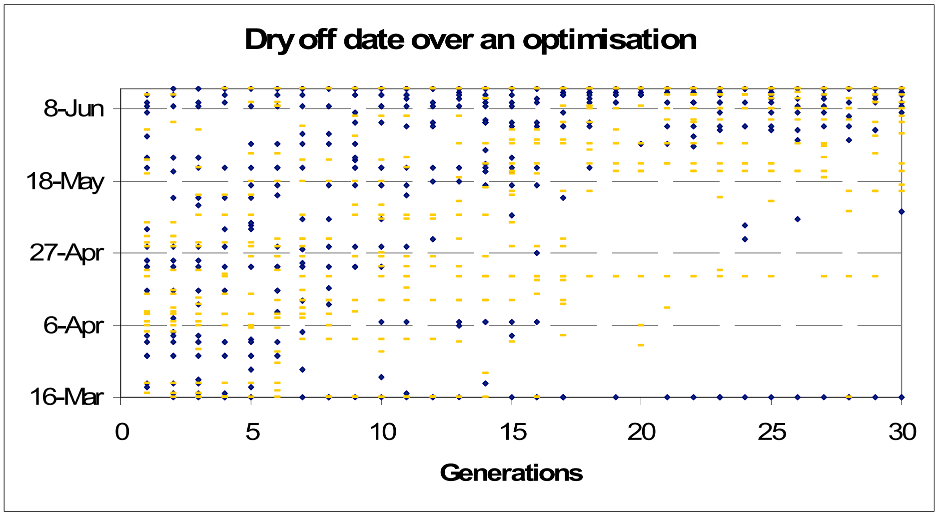

Figure A.5:

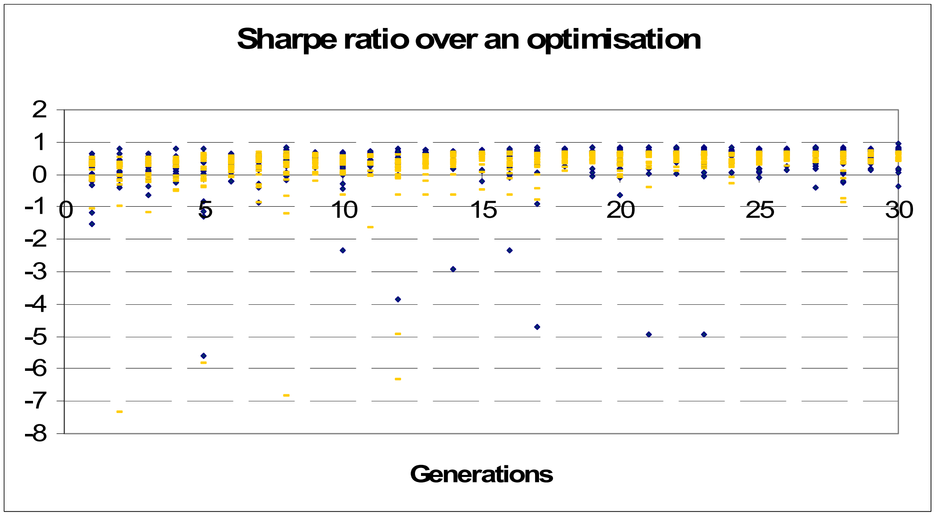

Figure A.6:

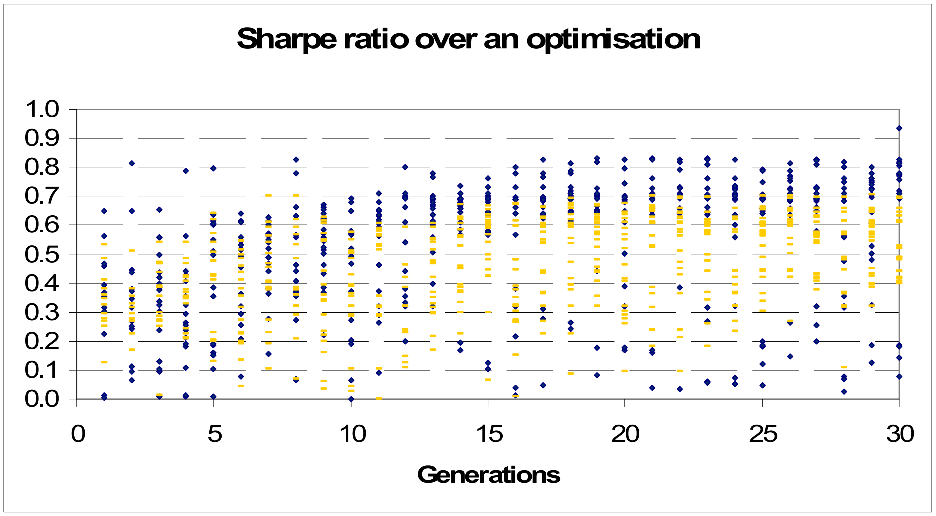

Figure A.7

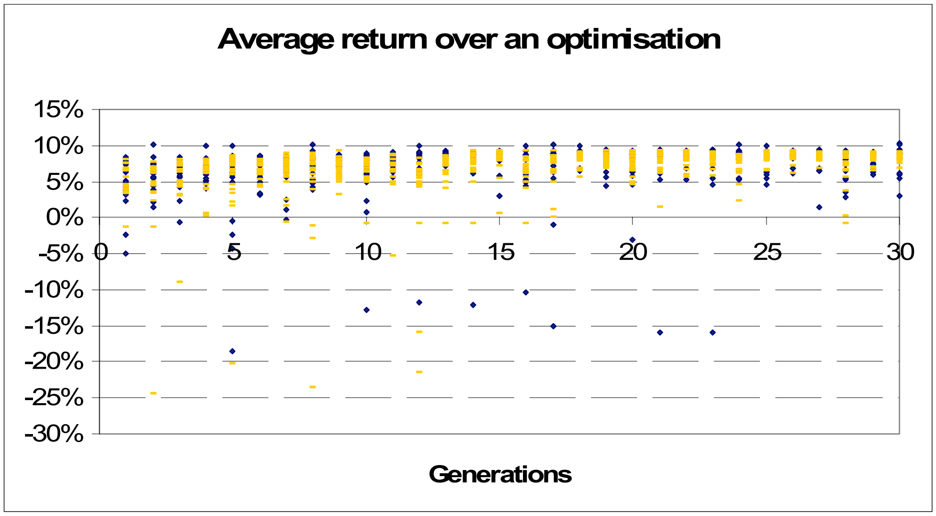

Figure A.8:

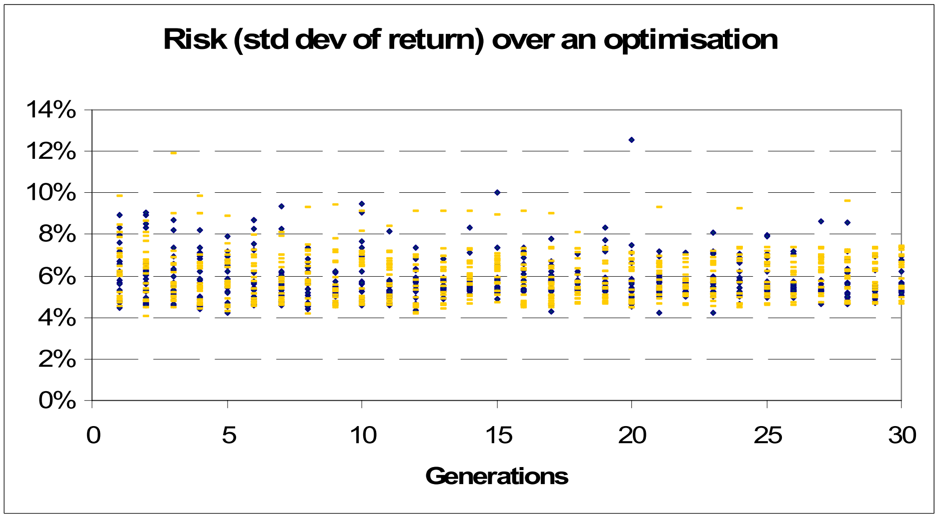

Appendix B: Parallel visualisation of all individuals.

Figure B.1: All results from an optimisation using Sharpe ratio as the fitness function

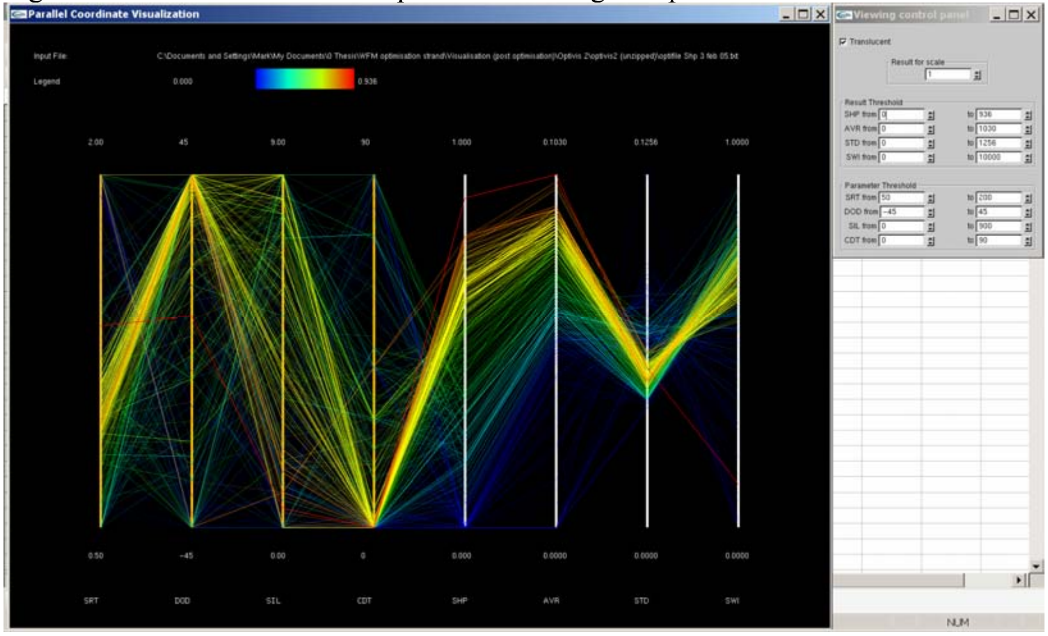

Input File: C:\Documents and Settings\Mark\My Documents\0 Thesis\WFM optimisation strand\Visualisation (post optimisation)\Optivis 2\optivis2 (unzipped)\optfile Shp 3 feb 05.txt

Legend 0.000 0.936

☑ Translucent

Result for scale: 1

Result Threshold
SHP from 0 to 936
AVR from 0 to 1030
STD from 0 to 1256
SWI from 0 to 10000

Parameter Threshold
SRT from 50 to 200
DOD from -45 to 45
SIL from 0 to 900
CDT from 0 to 90

Figure B.2: All results from an optimisation using Stochastic Dominance

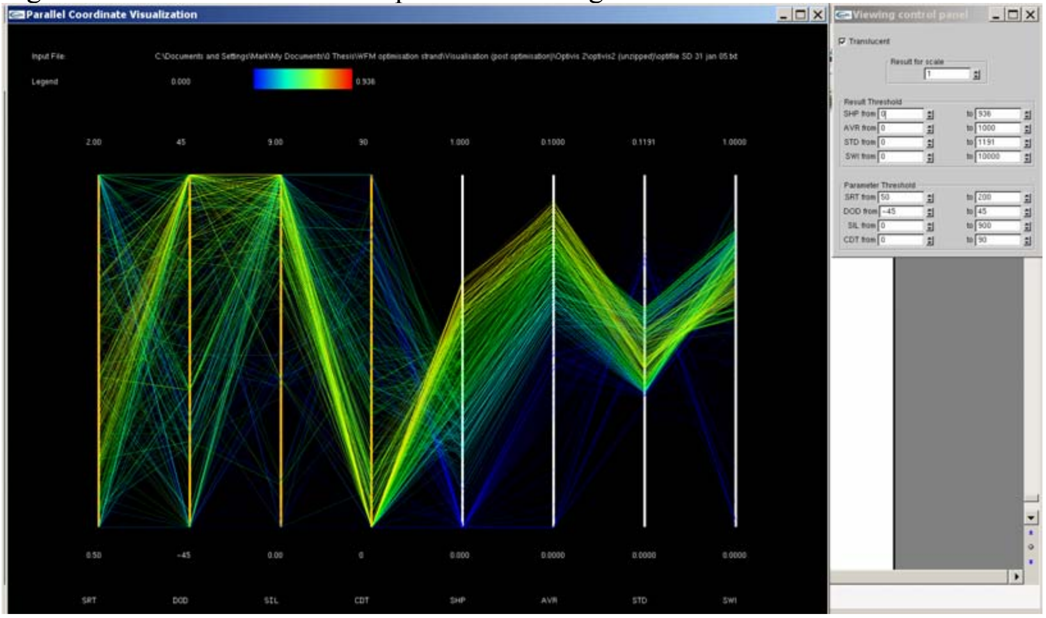

Input File: C:\Documents and Settings\Mark\My Documents\0 Thesis\WFM optimisation strand\Visualisation (post optimisation)\Optivis 2\optivis2 (unzipped)\optfile SD 31 jan 05.txt

Legend 0.000 0.936

☑ Translucent

Result for scale: 1

Result Threshold
SHP from 0 to 936
AVR from 0 to 1000
STD from 0 to 1191
SWI from 0 to 10000

Parameter Threshold
SRT from 50 to 200
DOD from -45 to 45
SIL from 0 to 900
CDT from 0 to 90

Four inputs

1. SRT: Stocking Rate, bounded between 1.5 (0.50) and 6 (2.00) cows per hectare
2. DOD: Dry Off Date, bounded between 16 March (-45) and 14 June (45)
3. SIL: Initial Amount of Silage, bounded between zero (0.0) and 4.5 (9.0) wet tonnes per cow
4. CDT: Calving Date, bounded between 2 June (0) and 31 August (90)

Four results

1. Sharpe ratio (0.0 to 1.0)
2. Average return (0 to 10.3%)
3. Standard deviation of returns (0 to 10%)
4. Shapiro Wilks p-value. (0.00 to 1.00)