Applying Differential Evolution to a Whole-Farm Model to Assist Optimal Strategic Decision Making

[1]**<u>Neal, M.</u>** [2]**Wastney, M.,** [2]**Levy, G.,** [1]**Drynan, R.,** [1]**Fulkerson, W., Post, E.,** [2]**Thorrold, B.,** [2]**Palliser, C.,** [2]**Beukes, P., and** [2]**Folkers, C.**

[1]Sydney University, [2]Dexcel, New Zealand, [3]Lincoln University E-Mail: mbneal@tpg.com.au

Keywords: *Optimisation; Dairy; Evolutionary algorithm; Differential evolution; Strategic decisions.*

EXTENDED ABSTRACT

A whole-farm model of a dairy farm was optimised to assist the strategic decision making given the uncertain environment. Decisions under consideration included the stocking rate, the calving pattern (defined by calving and dry-off dates) and the use of supplementary feed. These individual decisions are considered as the choice of a farm system in this paper.

The farmer's decision making problem is simplified by assuming a perfect labour market, and so the farmer will be primarily concerned with the return on equity. Farm systems are compared with stochastic dominance or the Sharpe ratio, depending on whether there is a perfect capital market and returns are normally distributed.

The Dexcel whole-farm model was extended with an economics component that allowed the calculation of farm returns. The production risk, due to different seasonal conditions, and the price risk from feed prices, milk prices and capital appreciation was also modelled.

The whole-farm model was optimised with the use of the differential evolution algorithm of Storn and Price (1997). The optimisation was parallelised over the Dexcel computer network allowing for results to be found 30-40 times faster compared to using a single computer.

Improved farm systems could be found more quickly by using the Sharpe ratio as a selection tool compared to stochastic dominance. This is despite the possibility that some assumptions for the use of the Sharpe ratio were not met (e.g., normality and perfect capital markets).

The inputs and results were each in four dimensions, making traditional visualisation difficult. Parallel visualisation was used to compare near optimal farm systems (figure 1). The optimal system (in red) is characterised by a long lactation length, a moderately high stocking rate and a low initial amount of silage.

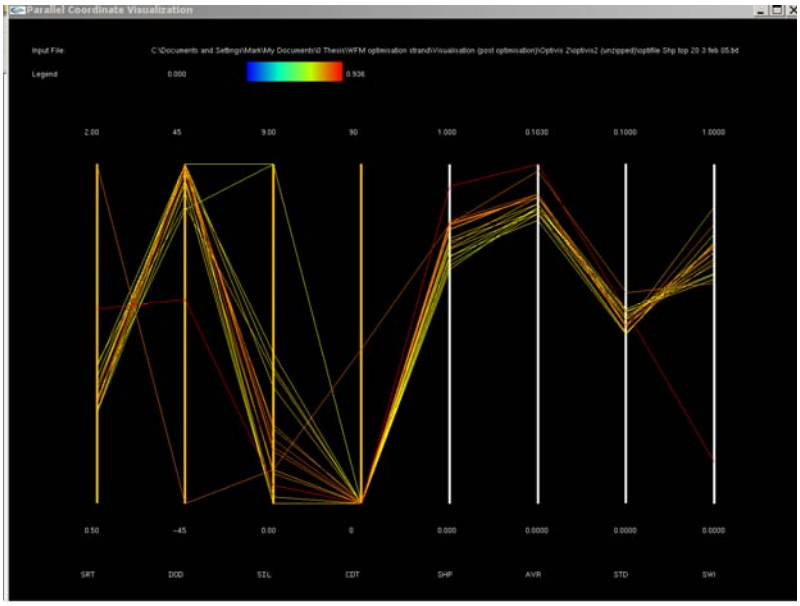

Figure 1. Near optimal farm systems, based on maximum Sharpe Ratio

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

1. INTRODUCTION

The central aim of this project was to improve dairy farmer's strategic (long term) decisions by finding optimal farm systems. The process involved the integration of an economics component into an existing whole-farm model, followed by utility based optimisation using a evolutionary algorithm. The optimisation was first conducted with a single objective function, the Sharpe ratio, assuming perfect capital and labour markets for a farmer with two investment alternatives, a dairy farm or a risk-free asset. The optimisation problem was expanded by relaxing the assumption of perfect capital markets and normally distributed returns. This prevents the use of the Sharpe ratio but allows stochastic dominance to be applied. Emphasis was placed on post-optimisation analysis, particularly through visual approaches.

2. DECISION-MAKING FOR DAIRY FARMERS

The particular focus of this paper is assisting the strategic decisions that the farmer must make, especially in reference to New Zealand dairy farms. A dairy farmer must make both strategic and tactical (short term) decisions when optimising returns. 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, how much supplementary feed to use, what speed to rotate paddocks and how much fertiliser to apply to each paddock over the next week.

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, combined together, describe a farm system.

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 aspects of climate including low or excessive rainfall, temperature and solar radiation impact on pasture growth.

The farmer's decision problem can be simplified somewhat by assuming that a farmer has flexibility in labour supply decisions. In other words, the farmer will hire labour if required, or sell labour if there is an excess of labour. Under this assumption the level of labour used by the farm system will not be important in deciding on a farm system. Hence the farmer's major concern will be with maximising the return on the other major input they supply to the farm, 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. A final assumption is that the farmer's utility function can be defined in terms of the mean and variance of returns, allowing the application of separation theorem. 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' risk preferences expressed through utility functions to indifference curves U_A and U_B. The risk free rate of borrowing and saving is shown as R_f. The 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, yielding higher utility (or satisfaction). 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, again achieving higher utility than would be possible without borrowing.

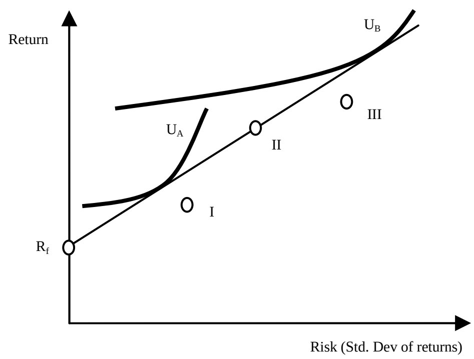

Figure 2. Two farmers choosing a farm system

The slope of the line from R_f is the Sharpe ratio (Bodie et al. 2005):

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

In other words, the Sharpe ratio is the excess return per unit of risk. The farm system that results in the maximum Sharpe ratio is the farm system that should be chosen by the farmer to optimise returns, 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 assumes only that the farmer prefers more to less.

While first degree stochastic dominance makes few assumptions, it may not be very discriminative. It may not significantly reduce the number of alternative farm systems that the farmer must choose the optimal farm system from. This is because a common result from the first degree stochastic dominance test is that neither system A nor system B dominates (or is clearly preferred to) the alternative.

Second degree stochastic dominance makes the additional assumption that some level of risk aversion exists. This rule (B dominates A) is summarised in Hardaker et al, (2004) and is more discriminatory than first order stochastic dominance.

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 reduced to 16 funds (78% of funds were dominated).

3. 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 with the primary tool used for modelling being 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 built in a different programming language. The WFM has a choice of two pasture models, 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 two animal models; currently a simple energetics based cow model (SimpleCow) and a more complex model (MollyCow) based on the work of Baldwin (1995) at the University of California, Davis.

Figure 3 shows a simplified schematic of the WFM (Wastney et al. 2002). 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 and 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. The biophysical output is then used in the economics component of the model to generate a simplified profit and loss statement, balance sheet and return on assets.

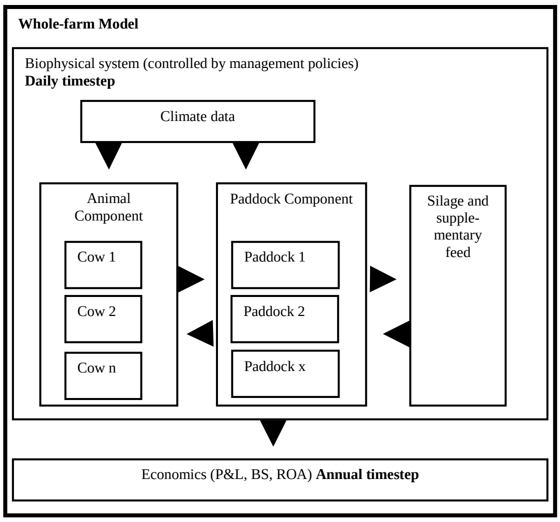

Figure 3. Whole-farm model schematic

A large proportion of costs are defined in an activity based costing framework. Default values are generated through the use of economic survey data specific to a farm region. The main cost drivers are the number of cows calved and the effective farm area. With further improvements in the model, new cost drivers can be incorporated to reflect the wider alternatives of management. Constant returns to scale is assumed and Neal (2004a) found this to be a reasonable assumption over a wide range of farm sizes encompassing more than half of the farms surveyed in the 2002 economic survey (Dexcel, 2003).

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:

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. the risk of land prices moving up or down.

The risk report takes a farm system and performs Monte Carlo simulation to find the distribution of returns with variable weather, milk price, supplementary feed price and capital appreciation rates.

4. OPTIMISATION WITH AN EVOLUTIONARY ALGORITHM

The evolutionary algorithm (EA) used for the optimisation of the Dexcel WFM is a variant of common genetic algorithms. This differential evolution (DE) 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. This is implemented in a way that benefits from adaptive search. Adaptive search implies that knowledge of the diversity of the current population is used in creating new individuals. In terms of the DE, a population is a group of individuals, where an individual is a specific farm system.

Firstly a member of the population is chosen as a Parent (P). A child (C) is then created and replaces the parent if it has a higher fitness level. To create the child three members are selected at random from the population. 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 4. Mayer et al (2005) also implements a useful method suggested by Kinghorn (pers comm.) to prevent premature convergence due to the interpolative nature of using an f value of less than one. The method is to "pulse' f to a value (much) larger than one, every *n* generations, facilitating extrapolative search.

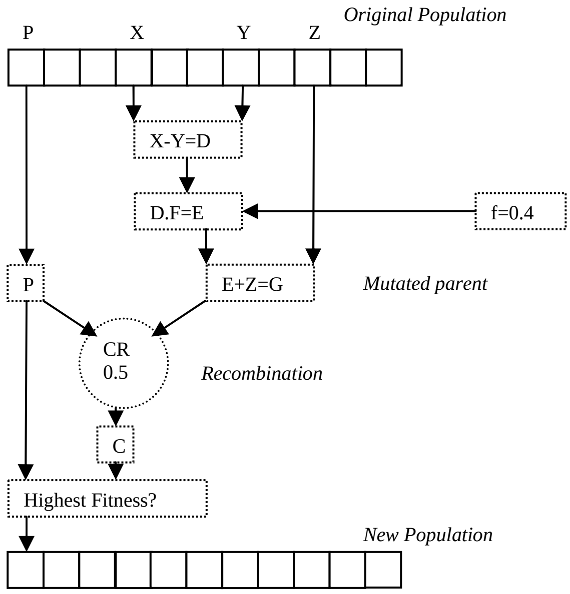

Figure 4. Representation of the differential evolution process; adapted from Storn (1997)

In the WFM selection was driven by a fitness measure which was implemented in two alternative ways. Firstly, through the maximum Sharpe ratio. This rule is very discriminating, as explained in Section 2. Secondly, through first order stochastic dominance. A common outcome of a stochastic dominance comparison is that no individual is clearly preferred to the other. 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. The calculation of the individual's fitness (running a WFM simulation and risk report), is typically computationally expensive, but can be calculated independently of the selection, recombination and mutation process of the EA. This feature is used to maximum advantage in the implementation of the WFM optimisation 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.

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. The communication between slave and master is via generic internet protocols.

5. RESULTS

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 four inputs to find the optimal farm system(s):

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 fi=0.4 and pulsed to fe=4.0 every 4 generations for extrapolative search. The Sharpe Ratio was calculated based on a risk free rate (rf) 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. 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. This is mainly due to its improved discriminatory power (figure 5). In the stochastic dominance optimisation, around 50% of the replacements were by dominance relationships, reflecting its reduced discriminatory power. Interestingly, the average Sharpe ratio of the children of each generation does not appear to be radically different using different fitness functions (figure 5). This would be due to the nature of the WFM returns and could not be expected to occur in all related problems.

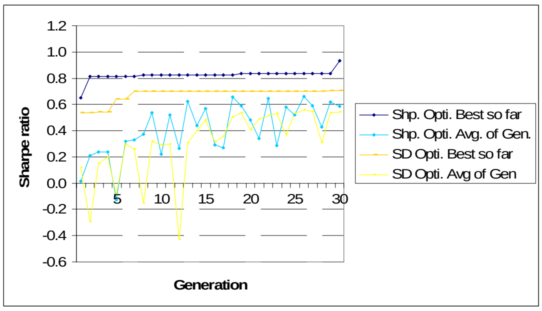

Figure 5. Sharpe ratio across an optimisation with two fitness measures

All individuals (farm systems) evaluated from both optimisations can be examined in the mean-standard deviation space (figure 6). 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 Sharpe ratio optimisation, although only during the last generation for the replicate shown. 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, however this is well below the highest Sharpe ratio found.

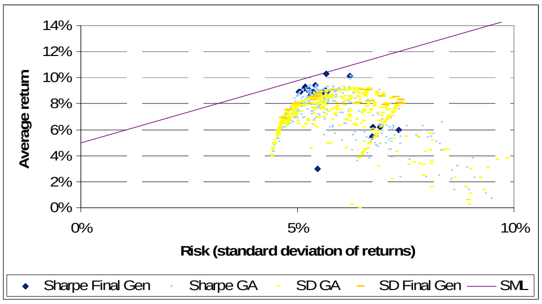

Figure 6. Individuals from optimisations in a mean-standard deviation space

The optimisation showed a tendency towards individuals whose cumulative distribution of returns are clustered closely and crossover in at least one point. This slows the progress of a stochastic dominance based optimisation.

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 (Shapiro et al. 1965). The null hypothesis is that the returns are normally distributed. 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. Hardaker (pers comm.) also advises that a statistical test for normality may not be sufficient for the assumption of normally distributed returns because a deviation from normality at low return levels has a higher utility weight compared with deviations at high returns.

Analysing 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 and 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, representing a very different type of farm system. The final generation is displayed in figure 1 using OptVis.

The final generation of the stochastic dominance based optimisation provided solutions that were quite different farm systems to those found 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 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.

Stochastic optimisations by their nature will not produce the same results each time they are repeated. For comparative purposes, 3 replicates were performed and their progress in the Sharpe ratio was recorded. 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 increasing the number of generations, increasing the pulse size and/or frequency of extrapolative search (f_e) or increasing the population size.

6. CONCLUSIONS AND FUTURE WORK

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.

Other basic limitations of the analysis 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 (treading damage) 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 be expected to positively bias towards high stocking rates and the high use of supplementary feed.

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.

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.

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 (2004b), 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.

Improvements to the WFM are incorporated regularly with the most recent version available can be used for optimisation. The short term direction of the current research is towards improved optimisation configuration and directly modelling the capital required for different feeding levels. Alternative fitness functions such as second degree stochastic dominance and certainty equivalent maximisation is also under analysis.

Longer term research may involve multi-objective optimisation which would allow preferences to be expressed 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.

In summary, there is a large potential for multidisciplinary work through a modelling project such as the WFM to link biophysical modelling to economic optimisation to assist the decision maker in making better strategic and tactical decisions.

7. ACKNOWLEDGMENTS

Mark's PhD was partly funded by a Dairy Australia scholarship. The cooperation of Dexcel was invaluable to the project.

8. REFERENCES

Baldwin, R.L. (1995), *Modeling Ruminant Digestion and Metabolism*. Chapman and Hall, London, UK.

Dexcel Ltd (2005) Dexcel – Research – Farm systems, [Online], Available from: URL http://www.dexcel.co.nz/main.cfm?id=248&sid=190&id2=199 [Accessed 13 September 2005]

Dexcel (2003) *Economic Survey of New Zealand Dairyfarmers 2001-2002*, Dexcel Ltd, Hamilton

Hardaker, J. B., Huirne, R. B. M. and Anderson, J. R. (2004), *Coping with risk in agriculture*, CAB International, Wallingford.

Levy, H. (1992) Stochastic Dominance and Expected Utility: Survey and Analysis *Management Science*, 38(4), 555-593

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

McCall D.G., and Bishop-Hurley G.J. (2003) A pasture growth model for use in a whole-farm dairy production model, *Agricultural Systems*, 76 1183–1205

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

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

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

Storn, R. and Price, K. (1997) Differential Evolution – A Simple and Efficient Heuristic for Global Optimization over Continuous Spaces, *Journal of Global Optimization*, 11: 341–359

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.