Farm Optimisation Group

Farm Optimiser Documentation

Contents

Farm Optimiser Overview

In summary, Farm Optimiser is a Python based, whole farm LP model. Farm Optimiser leverages a powerful algebraic modelling add-on package called Pyomo (Hart et al., 2011) and a mathematical solver to efficiently build and solve the model. The model represents the economic and biological details of a farming system including components of rotations, crops, pastures, livestock, stubble, supplementary feeding, machinery, labour and finance. Furthermore, it includes land heterogeneity by considering enterprise rotations on any number of soil classes. The detail included in the modules facilitates evaluation of a large array of management strategies and tactics.

Farm Optimiser has been built with the aim of maximising flexibility. Accordingly, depending on the problem being examined, the user has the capacity to:

- Change the region or property.
- Select the level of dynamic representation. For example, the user controls the number of discrete options for seasonal variation and price variation.
- Add or remove model components such as the number of land management units, land uses, novel feed sources such as salt land pasture, times of lambing for the flock and flock types (pure bred, 1st cross or 2nd cross).
- Adjust the detail in linearising the production functions (e.g. the number of livestock nutrition profiles).
- Make temporary changes to production parameters and relationships. For example, altering the impact of livestock condition at joining on reproductive rate.
- Constrain management. For example, fix the stocking rate or crop area.
- Include or exclude farmer risk aversion.

To facilitate user flexibility and support future development, Farm Optimiser is built in Python, a popular open source programming language. Python was chosen over a more typical algebraic modelling language (AML) such as GAMS or Matlab for several reasons. Firstly, Python is open source and widely documented making it easier to access and learn. Secondly, Python is a general-purpose programming language with over 200,000 available packages with a wide range of functionality (Van Rossum, 2007). Packages such as NumPy and Pandas (McKinney, 2012) provide powerful methods for data manipulation and analysis, highly useful in constructing Farm Optimiser which contains large multi-dimensional arrays. Packages such as Multiprocessing (Singh et al., 2013) provide the ability to run the model over multiple processors taking advantage of the full computational power of computers to significantly reduce the execution time of the model. Thirdly, Python supports a package called Pyomo which provides a platform for specifying optimization models that embody the central ideas found in modern AMLs (Hart et al., 2011). Python's clean syntax enables Pyomo to express mathematical concepts in an intuitive and concise manner. Furthermore, Python's expressive programming environment can be used to formulate complex models and to define high-level solvers that customize the execution of high-performance optimization libraries. Python provides extensive scripting capabilities, allowing users to analyse Pyomo models and solutions, leveraging Python's rich set of third-party libraries designed with an emphasis on usability and readability (Hart et al., 2017).

The core units of Farm Optimiser are:

1. Inputs: The model inputs are stored in a multi-layered database for quick access in the app.

2. Precalcs: The precalcs are the calculations applied to the input data to generate the data for the Pyomo parameters (in other terms, the conversion of the inputs to the parameters for the LP matrix). The precalcs for each individual trial (trial is the name for a single model solution) can be controlled by the user with the 'experiment' spreadsheet which allows inputs from the three input spreadsheets to be temporarily adjusted, or the intermediate calculations in the precalcs to be temporarily adjusted.

3. Pyomo and solver: This is the LP component of the model (matrix generation). It defines all the decision variables, the objective function, the constraints and parameters then utilises them to construct the model's equations (i.e. constraints). Components of the LP model can also be temporarily adjusted by the user via the 'experiment' spreadsheet. Pyomo formulates all the equations into a linear program format and passes the file to a solver.

Simply put, the LP component of Farm Optimiser sees a range of activities (variables) that can be selected so long as all the constraints are met. In many cases this is easy to visualise for example Farm Optimiser can select as many or as few labour sources as it likes so long as all on-farm tasks can be completed by a suitably skilled staff member. However, in some cases it is a little more complex. For example, in the pasture module where an activity represents a current state (e.g. FOO at the start of the period 1), a management option (e.g. level of grazing during period 1) and a future state (e.g. FOO at the end of the period 1). The pasture activities in period 1 are then transferred to period 2 using constraints that ensure that the FOO at the start of period 2 is equal to the FOO at the end of period 1. The calculations that determine the future states of an activity based on a starting point and a range of management options are conducted in the precalcs. These are similar to that of simulation calculations.

The procedure for building and solving Farm Optimiser is that firstly, the inputs are read in from the Excel files. The experiment spreadsheet is read that includes the temporary adjustments (sensitivities) for the model parameters. Furthermore, the spreadsheet allows the user to group trials into an experiment to be run as a batch. For example, the user may be interested in the impact of increasing prices, hence an experiment examines several price levels. Secondly, each module containing precalcs is executed. The parameters produced are stored in a python data structure called a dictionary. Then the Pyomo section of the model creates the decision variables, formulates the model constraints, populates the parameters with the coefficients from the precalcs and passes the information to a linear solver. The results from the solver reveal the maximum farm profit and the optimal levels of each decision variable that maximises the farm profit (or some other objective function). From here the user can create a range of reports.

Key improvements

Some of the key improvements of Farm Optimiser over previous optimisation models include;

i. Inclusion of price and weather uncertainty, the associated short-term management tactics and farmer risk attitude.
ii. Increased rotation options.
iii. Extra detail on the biology of livestock production that allows:
    a. Inclusion of optimisation of the nutrition profile of livestock during the year.
    b. A larger array of livestock management options such as time of lambing and time of sale.
iv. Improved pasture representation that includes production effects of varying grazing intensity.
v. More detailed representation of crop residue that includes multiple feed pools based on quality and quantity.

Additionally, developing Farm Optimiser in Python has resulted in a flexible framework that overcomes many previous structural challenges such as scalability. The structure allows the user to alter the biological detail to balance computer resource requirement against model realism in different aspects of the farm system. For example, the user of Farm Optimiser can easily alter the number of discrete options represented in different sections of the model so that detail can be added to aspects that are important for a particular analysis while simplifying the less important. Furthermore, Farm Optimisers usability and detailed representation of the farm system means it can be applied to a plethora of current and future farming system opportunities and problems.

Rotations

Modelling of cropping or crop-pasture rotations to date has primarily been based on a predetermined restricted set of rotations represented as “activities” in a LP matrix (Wimalasuriya and Eigenraam, 2000). However, this approach often limits the potential rotations to be selected and does not capture the flexible nature of real-life rotation selection especially in the face of unfolding seasonal conditions. For example, using a fixed rotation structure, it is not possible to alter the rotation in response to the timing of early season rainfall. It also results in the necessity to build entirely new modules for each agro-climatic region due to differences in crop and rotation choices that are available and applicable to each region.

In Farm Optimiser, we adopt an alternative method proposed by Wimalasuriya and Eigenraam (Wimalasuriya and Eigenraam, 2000) , where the “activities” in the model are rotation phases. A rotation phase is a land use with a specific sequence of prior land uses (‘history required’). A constraint is included to ensure that for the model to select a given rotation phase, the ‘history required’ must match the ‘history provided’ from another rotation phase. The model solves for the optimal rotation through a selection of rotation phases. This is an unrestricted approach that supports a large range of possible rotations and allows greater flexibility for adding new land uses. Additionally, the approach aligns closer to reality, facilitating a more detailed and accurate representation of the effects of weather-year type on rotation choice.

Each rotation phase requires a history and provides a history. As a simple example, consider the rotation phase barley – wheat: canola in which canola is the current land use. Barley followed by wheat is the history required and wheat followed by canola is the history provided. Based on the current land use and the land use history the level of production (grain and stubble production from crop phases and, seed set and germination from pasture phases), the costs, the machinery requirement and the labour requirement are determined.

The rotation phases are designed to be as simple and general as possible while still capturing important performance and management variants. The system employed is to generate all possible combinations of the land use sequences over a set number of years, then the infeasible options are removed and the remainder are generalised where possible e.g. wheat, barley and oats may be generalised to cereal in the phase history if the type of crop does not affect subsequent productivity or costs.

The length of the rotation phases and the level of generalisation is determined so that the impacts of the history on the current land use production and costs are captured. These can be summarised by:

1. The need to track the number of crop phases to determine if an annual pasture needs reseeding.
2. The need to track the effect of a land use on the productivity or costs of subsequent land uses. This can be either:
    a. Fixing of soil nitrogen and its subsequent effect on following crops. This requires tracking:
        - The number of years of the legume as it affects the quantity of organic nitrogen.
        - The number of years since the legume to determine the remaining nitrogen.
    b. Impacts on disease levels.

    c. Impact on weed seed levels.

3. The impact of cropping on subsequent annual pasture seed bank and germination.

The impacts and assumptions of land use history on production and costs that are being captured in the rotation phases developed are:

1. Annual pasture will be resown if the four most recent land uses in the history are crops. Resowing impacts the current year and the succeeding year.
2. Lucerne (or Tedera) will be resown if the immediately preceding land use is not Lucerne (or Tedera).
3. The impacts of spray-topping and manipulating pastures lasts for two years.
4. Germination of annual pasture is affected by:
    a. The two most recent land uses in the history.
    b. The crop type immediately prior to the annual pasture. Specifically:
        - Germination is higher after an oat fodder crop.
        - A pulse crop increases growth of annual pastures (which is represented by an increase in germination).
5. A history of legume pasture (annual, Lucerne and Tedera) provides organic nitrogen for subsequent non-legume crops (cereal or Canola).
    a. The amount of organic nitrogen increases up to four years of consecutive legume pasture.
    b. The impact of the organic nitrogen lasts for a maximum of three years.
6. Pulse crops provide organic nitrogen for subsequent non-legume crops.
    a. The impact of the organic nitrogen lasts for a maximum of three years.
7. Leaf disease and root disease builds up for each land use and reduces productivity for consecutive land uses.
    a. It is assumed that the maximum level of disease is reached after 4 consecutive years of a land use.

b. There is variation in the length of the break (interval in years between the same land use) required and the duration of the benefits of a break.

To capture all the factors listed above, the length of the rotation phases represented in Farm Optimiser is defined to five years, allowing a history of four pastures to be tracked. To reduce the number of rotation phases, land uses in the history that are assumed to have the same impact on the production and cost of the current land use are grouped into 'land use sets' (see the online documentation).

Some of the rotation phases constructed will be illogical and are removed. For example, annual pasture is only resown after four years of continuous crop therefore any rotation phases that are generated with resown annual pasture that do not have four years of crops preceding it can be removed. To further reduce the possible number of rotation phases in the model, unprofitable and unused land sequences are removed. See RotGeneration for the full list of rules.

Cropping

Cropping is often a large component of broadacre farming in Western Australia (Planfarm, 2022). Crops are primarily established with a goal of harvesting and selling the grain for human or animal consumption. However, crops can also be used as fodder for farm livestock. Using a range of user inputs, the crop module calculates the crop production, the associated costs, the labour requirements and the machinery requirements. The inputs include crop yield, fertiliser and chemical requirements, frost damage, seeding rates, soil type, machinery size, paddock efficiency of seeding and harvesting, proportion of helper labour required, crop residue management, proportion of arable area, commodity prices, fees and levies. To accurately reflect rotation history, soil type and weather effects on crop production and management, many of these user inputs vary by rotation phase history, soil type and weather conditions. For example, the rotation phase history can influence the soil fertility, weed burden and disease and pest levels. These factors impact the potential yield and the level of fertiliser and chemicals.

Each phase provides a certain amount of biomass based on the inputs above. Accordingly, Farm Optimiser can optimise the area of each rotation phase on each LMU and the best way to utilise the biomass of each rotation phase. Biomass can be either harvested for grain, baled for hay or grazed as standing fodder. Farm Optimiser does not currently simulate the biology of crop plant growth under different technical management. Thus, Farm Optimiser does not optimise technical aspects of cropping such as timing and level of fertiliser applications. However, the user has the capacity to do this manually by altering the inputs (à la management of inputs in simulation modelling) or by including additional land uses which represent varying levels of inputs and production.

There are two methods that can be used to generate cropping inputs for the model:

1. Manually enter the inputs for selected rotation phases: The user can manually input the fertiliser and chemical requirements and resulting yield of each rotation phase. To do this accurately requires an in-depth knowledge of cropping in the location being modelled. Thus, the process is often done in collaboration with a consultant or specialist in the field. This input method can be limiting if the user is hoping to include a large number of rotation phases or land uses that are not well established in the given location because it can be difficult to determine accurate inputs.

2. Generate using simulation modelling (e.g. Thamo et al., 2017): APSIM is a whole farm simulation model widely used in Australia. APSIM has detailed modules which use robust relationships to simulate plant growth. The parameters used in APSIM can be altered to represent plant growth in many different situations. For example, different rainfall or soil conditions. Farm Optimiser users can use APSIM to generate yield of crops in a given rotation under a specified fertiliser and chemical regime.

Although both methods can be sources of yield estimates; only the first method is a source of estimates of required fertiliser and chemical at a particular location.

The crop management decisions that are optimised can include:

1. Area of each rotation phase on each soil type in each weather year depending on paddock history.
2. Area of each crop harvested, baled or grazed.
3. Contractual services for seeding or harvesting.
4. Labour allocation.
5. Time of sowing.

The model can also represent and compare (but not optimise in a single model solution):

1. Fertiliser application rate and timing.
2. Chemical application rate and timing.
3. Seeding rate.
4. Alternative cultivars.
5. Seeding technology.

Feed budget

Energy is the primary nutritional constraint for extensive ruminant livestock enterprises (Rickards and Passmore, 1977). As such, energy is the only nutritional element that is constrained in Farm Optimiser to ensure that feed supply is greater than or equal to the feed demand of the livestock (as measured by metabolisable energy). There is also a volume constraint that limits the minimum diet quality to ensure that the voluntary feed intake capacities of the livestock are sufficient to consume the quantity of feed selected. The volume of each feed source (kg of intake capacity / kg of feed dry matter) varies depending on the feed quality (relative ingestibility) and feed availability (relative availability) using relationships from Freer et al. (2007).

The feed supply from pastures, crop residues and supplementary feeds, is represented by changes in the type, amount and quality of feed available during the year. The feed demand of livestock is represented as the requirement for metabolisable energy and the feed intake capacity. The year is partitioned into 10 feed periods. A feed budget is carried out for each feed period to ensure that the feed demand of the flock can be met from the feed available on the farm. The dates of the feed periods during the growing season are selected to group periods that have similar supply and demand characteristics. During the growing season this is driven by the response of pasture growth to defoliation and the periods are shorter after the break of season and just prior to senescence. During the dry feed phase the dates are selected to minimise feed variation within each period and are shorter after pasture senescence and prior to the break of season. The selection of the period definitions is likely to alter depending on the region being modelled.

Any of the 10 periods can be the period that limits the farm carrying capacity. This is representing far more detail than underpins a typical gross margin analysis that considers a pre-defined feed limiting period of the year. Furthermore, Farm Optimiser includes the capacity to alter the live weight (LW) profile and hence feed demand of any class of stock in any feed period with a concomitant change in production per head. This links to the capacity for supplementary feeding the livestock to optimise the number of livestock carried on the farm. As such Farm Optimiser is much more detailed than a typical gross margins analysis of livestock profitability. If Farm Optimiser is compared to a simulation model the feed periods are equivalent to the time-steps in the simulation model, however, they are much longer than a typical simulation model that often considers daily time-steps. As such Farm Optimiser represents the feeding options in less detail than is possible in a dynamic simulation model, however, Farm Optimiser has the advantage of optimising the grazing management of the pastures and crop residues and optimising the target nutrition for each class of stock during the year.

Cross subsidisation of volume is a problem that can occur in the feed budgets of linear programming models. Cross subsidisation occurs if animals with divergent quality requirements are constrained by single energy and volume constraints; the single constraint is termed a feed pool. For example, consider two animals, one losing 100 g/hd/d and one gaining 150 g/hd/d. The first animal can achieve its target on low quality feed whereas the second animal needs high quality feed. However, if both of these animals were constrained using a single feed pool, then the total energy requirement and total intake capacity is combined, such that feeding medium quality feed to both animals meets the constraints. This is likely to be the optimal solution because the cost of feed by quality is a convex function and therefore the cost-minimising solution is to provide an average quality to both classes of stock. However, this is not a technically feasible solution. To reduce the possibility of cross-subsidisation of volume while still limiting model size, the energy requirement and maximum volume constraints are applied in multiple nutritive value pools, each spanning a small range of nutritive value (where nutritive value = ME requirement / volume capacity). This is more efficient in reducing model size and complexity than having a feed pool for each animal class.

Feed supply

The main sources of feed considered in the model are; pasture (annual and/or perennial), crop residue (stubble) and supplement (grain concentrates and conserved fodder). Farm Optimiser also includes some novel feed sources such as early season crop grazing, grazing standing fodder crops and salt land pastures.

The feed management decisions that are optimised can include:

1. Area of each pasture variety on each soil type.
2. Area of reseeded pasture based on paddock history.
3. Area of pasture manipulated and/or spray-topped based on paddock history and setting up for future land uses.
4. Grazing intensity of different pasture varieties on different soil types at different times of the year which manifests as a FOO profile of the pasture.
5. Timing and extent of pasture deferment.
6. Level and timing of supplementary feeding of hay or grain to each class of stock.
7. Grazing management of stubbles.
    a. The time to start grazing of each stubble.
    b. The class of stock that grazes the stubble.
    c. The duration of grazing.

    d. The amount of supplementary required in addition to stubble (to meet alternative LW profiles).
8. Area of fodder crops established and their grazing management.
9. Tactical grazing of standing crops in place of harvesting.
10. Amount of early season crop grazing.
11. Salt land pasture grazing management.
12. Conserving surplus pasture as hay or silage.
13. The level of growth modifier (e.g. nitrogen fertiliser) applied to pasture.

The model can also represent and compare (but not optimise in a single model solution):

1. The level of phosphate fertiliser application to pastures.
2. The impact of varying pasture conservation limits.
3. Altering pasture cultivars on different land management units.

Pasture

Pasture is the primary livestock feed source because in an extensive farming system it is a cost-effective source of energy available for the entire year. Different pasture types can be represented by altering the inputs for each pasture type. The default pasture type is “annual pasture”. However, by altering the inputs, perennial pastures and mixed swards can be represented.

Pastures are often included in a rotation to provide a break from cropping, which can rejuvenate soil conditions, provide disease and pest management and provide a cheap feed source for livestock. The pasture module generates the pasture production (as discussed below) and the costs and labour associated with seeding, monitoring, fertilising and spraying. Similar to cropping, the inputs vary based on rotation history, soil type and weather. Farm Optimiser can then optimise the area of each pasture phase to include on each land management unit (LMU).

Pasture feed sources can be supplemented with concentrates, and in a mixed crop-livestock farm system the pasture can be complemented with dry residues from crop production (stubbles). The biology and logistics of pasture growth rate that are represented in Farm Optimiser is:

- Pasture growth rate (PGR) is dependent on pasture leaf area, which is quantified by the level of feed on offer (FOO, kg of DM/ha). Additionally, PGR for each pasture type varies with the phase during its life cycle, soil moisture, sunlight and level of growth modifier applied. All are quantified by their land management unit (LMU), time of year and weather-year.

- The average FOO during a feed period depends on FOO at the beginning of the period, the grazing intensity and the PGR during the period.

- The mobilisation of below-ground reserves (germination) of annual pastures at the break of season is dependent on the seed bank. The seed bank is controlled by the rotation in which the pasture is grown and varies with LMU.
- The mobilisation of below ground reserves of perennial pastures at the break of season can also be adjusted by rotation. However, perennials usually are not grown in rotation with crops.
- The maximum intake of animals grazing pasture depends on FOO and diet dry matter digestibility (DMD). Intake can be less than maximum which implies that the optimum solution can include rationing of animal intake via rotational grazing.
- The digestibility of the diet selected by animals grazing green pasture depends on the sward digestibility and the animal's capacity for selective grazing. Sward digestibility varies depending on the pasture species, the time of year, the LMU and the FOO of the pasture. Selectivity depends on FOO and grazing intensity.
- Dry pasture that is not consumed is deferred to later in the year, with a reduction in both its quality and quantity. Livestock can select a higher quality diet when first grazing the dry pasture but quality reduces with extra grazing.
- Livestock trample both green and dry pasture while foraging in proportion to the amount consumed.
- The risk of resource degradation increases when ground cover is lower so there is a user defined minimum limit to ground cover during both the green and dry phases of the year.

The decision variables optimised in Farm Optimiser, that represent the above biology are the:

- rotation phases in which pasture can be grown on each LMU.
- FOO profile during the year that is represented by a discrete range of FOO levels at the start of each feed period.
- grazing intensity and the variation across feed periods during the year is represented by a discrete range of the severity of defoliation in each feed period.
- level of growth modifiers (nitrogen or gibberellic acid) applied to the pasture.
- quantity of dry feed consumed from each of 2 dry feed quality groups in each feed period.

The nutritive value of pasture is determined by the metabolisable energy per unit of dry matter, the relative ingestibility and the relative availability. This varies with:

1. Feed period.
2. The level of FOO. The greater the FOO, the lower the average digestibility of the sward. Lower digestibility of a high FOO sward is associated with the lignification that occurs in older foliage. Higher digestibility of a low FOO sward is associated with the higher digestibility of new growth that constitutes a higher proportion of the sward. There can be some error associated with this assumption if the low FOO was generated by grazing a high FOO sward back to a low FOO, in which case most of the DM would be stalk at the base of the plant which compares to a sward maintained at a low FOO level since the break of the season.
3. Grazing intensity. With heavy grazing there is little scope for selection, so the diet digestibility equals the sward digestibility. With light grazing there is scope for selection and diet quality that approaches that of high quality leaf. Note, increasing the energy content of the feed also improves the ingestibility of the feed (Freer et al., 2007).

Pasture on non-arable areas in the crop paddocks is modelled as above with a few additions. Firstly, pasture on non-arable area is represented as a continuous annual pasture. Secondly, non-arable pasture on crop paddocks is not available for grazing until after harvest and therefore it goes into the low-quality dry feed pool. Accordingly, pasture on non-arable areas of the crop paddocks does not receive any farm inputs.

Pasture grazed on the crop paddocks in the period before destocking for spraying and seeding is represented as a pre-specified quality and maximum quantity available each day on the area that is yet to be seeded, with the additional requirement that pasture must be destocked 10 days prior to seeding to allow time for an effective knockdown spray.

Crop residue

At the end of the growing season Farm Optimiser has the option of harvesting or baling each crop, which leaves stubble for stock consumption, or crops can be left standing for fodder grazing. Stubble and fodder are modelled in the same ways, as follows. In general, sheep graze crop residues selectively, preferring the higher quality components. Thus, they tend to eat grain first, followed by leaf and finally stem. To allow the optimisation of the quantity of the stubble grazed, and to reflect selective grazing the total crop residues are divided into ten categories. The higher categories are better quality but generally lower quantity. Consumption of a higher quality category allows the consumption of a lower category (e.g. sheep cannot consume any of category B until some of category A has been consumed).

The total mass of crop residues at first grazing (harvest for stubble and an inputted date for fodder) is calculated as a product of the biomass, harvest index and proportion harvested. Over time if the feed is not consumed it deteriorates in quality and quantity due to adverse effects of weather and the impact of sheep trampling.

Residue production can be positively impacted by frost because frost during the plants flowering stage can damage cell tissue and reduce grain fill (Zheng et al., 2015). This results in less grain and more residue due to not using energy resources to fill grain. Thus, the harvest index used to calculate biomass to residue is adjusted by a frost factor.

Supplement

Supplementary feeding is the supply of additional feed to livestock, primarily grain and hay (which are both represented in the model). Supplementary feeding is commonly used to help meet production targets such as lamb growth rates prior to sale, or to fill the feed gap to allow higher stocking rates during the summer and autumn months when pastures and crop residues are limiting. Additionally, feeding supplements can be used as a tactic to allow pastures to be deferred early in the growing season which increases subsequent pasture growth rates through increasing leaf area index.

Farm Optimiser represents a range of supplements including, oats, lupins and hay. Grain and hay as supplementary feeds can either be grown on farm or purchased from another farmer at a farm-gate price (i.e. net price of a product after selling costs have been subtracted) plus the transaction and transport costs. Supplementary feeding incurs a depreciation cost associated with storage infrastructure and variable costs associated with insurance, silo preparation, insect management, grain shrinkage/losses and machinery usage when feeding the supplement. The costs are calculated per tonne to allow for variations in grain density and amounts fed. Supplementary feeding also incurs a labour requirement for time spent traveling to and from the silo, filling the sheep feeder, emptying the feeder, and transporting between paddocks.

Crop grazing

Crop grazing is an option that allows stock to graze green crops, by default, from June until August (user customisable range). Green crops have a high energy content and grow erect allowing for easier grazing. Therefore, crops can meet livestock energy needs at a lower FOO than an equivalent pasture. However, for every kilogram of crop biomass consumed yield is reduced by 150 grams per hectare (user customisable), with a corresponding effect on stubble production.

Salt land pasture

Salt land pastures (SLP) are a novel feed source that consists of saltbushes and a grazable pasture understory. SLP establishment requires labour, specific machinery and a significant financial outlay, however it comes with numerous characteristics which make it attractive for certain situations. These characteristics include:

1. Saline tolerance and can therefore be established on land management units that would have had very low productivity (Barrett-Lennard et al., 2003).
2. Draw-down of the water table (Barrett-Lennard and Malcolm, 1999). This drawdown allows salts to be flushed from the topsoil of the moderately saline land, thereby creating growing conditions more suited to higher productivity annual pastures and perhaps leads to long term rehabilitation of the area.
3. Edible leaf for livestock consumption (O'Connell et al., 2006).
4. Livestock shelter – shelter provided by shrubs can be used by stock at vulnerable times such as lambing which helps increase animal survival.
5. Increased wool growth due to additional nutrients provided by grazing saltbush (Norman et al., 2010).
6. Reduced erosion risk due to the wind protection provided by the saltbushes year-round (Barrett-Lennard et al., 2003).

The salt land pasture land use is a combination of saltbush and understory. The saltbush module represents the saltbush component and the understory component is calculated in the pasture module.

The saltbush module includes the:

- cost of salt land pasture establishment and maintenance.
- productivity of saltbush during the year based on grazing management.
- feed value of saltbush during the year.
- diet selectivity of salt bush versus understory.
- impact of salt consumption on animal intake.

Livestock

A powerful and advanced feature of Farm Optimiser is its ability to optimise livestock liveweight/nutrition profiles. Farm Optimiser does this by generating production parameters for animals following a range of nutrition profiles (up to 2000 profiles for each class of sheep can be concurrently evaluated). These are represented as different decision variables which allows the optimisation of a wide range of management decisions. The total feed requirement and the minimum diet quality can vary for each feed period for each livestock decision variable. The range of nutrition levels are represented by profiles that are continuous for the entire year. At the end of the nutrition cycle (year) the range of final liveweights are 'condensed' back to a range of starting weights for the start of the next nutrition cycle. This capacity allows Farm Optimiser to differentially feed animals based on reproduction, sale goals and feed supply based on land use selection while minimising model size and computing resources required.

Farm Optimiser includes a livestock data generator that generates the production parameters for livestock for a user specified number of nutritional profiles. It is based on the relationships that underpin the GrazPlan suite of models as described by (Freer et al., 2007) and updated with production relationships developed in other research projects. Data is generated for the following components:

- Animal liveweight and sale values.
- Energy requirement profile and the nutritive value to achieve the target whole body energy profiles.
- Fleece production data; both quantity and quality (including the impact of the ewe liveweight profile during pregnancy on the lifetime performance of the progeny).
- Dam reproductive rate; represented by the proportion of ewes that are empty, single-, twin- and triplet-bearing.
- Perinatal survival of single, twin and triplet born lambs.
- Perinatal ewe survival associated with pregnancy toxaemia, dystocia and lambing difficulties.
- Mortality rates of dams, progeny and dry animals related to nutrition level.
- Foetal growth rates and birth weights for progeny.
- Milk production and progeny weaning weights.
- Husbandry cost and labour requirements.

- Methane emissions.

The values above are calculated for each ‘class’ of stock, for its lifetime. The feed supply offered to the animals is not based on simulating a growing pasture, rather the FOO, digestibility and supplement offered are inputs to the data generator. The outcome is the parameters required to define the production possibilities that are included in the matrix of the Farm Optimiser model.

The data generator model simulates the sires, dams and offspring from weaning to their latest possible sale age and simulates the young at foot from birth to weaning. The initial animal for the sires, dams and offspring are based on input values for liveweight, clean fleece weight and fibre diameter.

The prediction equations included in the data generator can be selected from a range of equation sources. Currently those source are:

1. GrazPlan equations as documented in Freer et al. (2012), which are an improved version of the Australian Feed Standards (SCA 1990).
2. Research trials carried out by Murdoch University, DPIRD and DPI Victoria that have quantified the impact of changing nutrition on production. This research began with the Lifetime Wool Trial (Oldham et al., 2011, Thompson et al., 2011) but has continued with a suite of other projects including Lifetime Maternals (Behrendt et al., 2019) and Mob size (Lockwood et al., 2019).
3. A selection of other sources that have developed equations to predict animal performance including:
    a. Blaxter and Clapperton for enteric methane emissions (Blaxter and Clapperton, 1965).
    b. NGGI for emissions including methane and nitrous oxide (DISER, 2021).
    c. Hutton Oddy’s group (NSW DPI) for alternative equations for heat production associated with maintenance and liveweight gain (Oddy et al., 2019).

In the full model the livestock management decisions that are optimised can include:

1. Number of animals carried (i.e. stocking rate) based on whole flock, whole year feed requirements and whole farm feed supply.
2. Sale age and weight of each animal group.

3. The proportion of the ewe flock mated to different sire genotypes (pure bred, maternal type or terminal)
4. The proportion of the ewe flock that is a first cross dam mated to a terminal sire (the dam cross is between the purebred and the maternal genotype)
5. The reproductive life of dams in the flock (based on whole flock feed requirements, value of wool variation by age, reproduction variation by age, the value of CFA dams at different ages, the selection pressure that can be applied on replacement ewes).
6. Whether to mate ewe lambs and the optimal proportion to mate.
7. A trading operation for dry animals. This can be either a short term trade with an aim to fatten animals or a multi-year trade to produce wool.
8. A ewe flock based on buying in ewes and mating all ewes to a terminal sire to produce first-cross lambs for sale. The age at purchase and sale can be optimised.
9. Diet selection for the animals based on the feed base options represented in the model including supplementary feeding.
10. Time in confinement and/or feed lot (note, this reduces the animals' energy requirements due to reduced walking).
11. Nutrition profile of the animals during the year which is related to reproduction status, wool value, sale objectives and unfolding climate conditions.
12. Differential feeding of dams based on litter size, lactation number and foetal age, provided the dams are pregnancy-scanned or assessed for 'gave birth and lost'
13. Optimal replacement policy based on:
    a. the change in reproduction and production over the animal's lifetime,
    b. the potential to increase per head production through culling and a response in the current generation.
14. Optimal weaning age for each dam age group.

Furthermore, constraints can be applied to the model to limit:

1. Level of enteric greenhouse gas emissions and emissions of nitrous oxide from faeces and urine.
2. Bare ground during the summer/autumn period
3. Animal mortality or liveweight loss during the feed limiting period of the year
4. Animals to graze at their voluntary feed intake level (i.e. that intake reflects the FOO and DMD offered to the animals i.e. feed is not rationed through active management of the stock. This has little effect unless the pasture management is also constrained to limit variation of the FOO and quality profile)

The model can also represent and compare (but not optimise in a single model solution):

1. The length of the joining period (measured in the number of cycles mated); including the trade-off between the number of ewes conceiving and the distribution of size and energy requirements of the later born progeny.
2. The age that the young ewes are mated. For example, a 7 month mating versus a 8.5 month mating for ewe lambs.
3. Accelerated lambing where ewes are mated every 8 months and therefore have 3 lambing opportunities in 2 years.
4. Variation in timing of lamb, hogget and adult shearing.
5. More frequent shearing. For example, adopting a shearing interval of 6, 8 or 12 months.

Finance

The financial components of the model include:

- interest
- cashflow
- a limit on capital borrowings
- minimum return on expenditure
- opportunity cost of assets

Each module tracks its relevant financial components.

To support the sporadic nature of farming income, finance is often drawn from the bank throughout the year to fund costly operations such as purchasing fertiliser and chemicals. In Farm Optimiser the total capital required for the given farm structure is tallied and can be constrained to a user specified level. This allows the user to examine how the business structure would change if finance is limited. This can also be used to ensure the model does not overdraw an unrealistic/undesired level of capital from the bank. Total farm capital required is calculated from the value of starting assets plus the sum of all the expenses minus any income, between the previous 'main' income (e.g. harvest or shearing) and the peak debt date. Peak debt is typically expected for an enterprise just before the main income is received for that enterprise, ensuring the main income for the enterprise is not included in the working capital constraint. The aim of the working capital constraint is to allow the user to constrain management practices which have high costs. If the main income was included in this constraint there would be no way to constrain high cost high reward management practices. The default is to have one peak debt date per enterprise, just before the main income for that enterprise is received.

In an equilibrium model there is no start and end point. This complicates the calculation of interest because interest must be calculated for a given period. In Farm Optimiser the interest period starts and finishes after the main income for the enterprise is received. This is logical from an expense point of view because the expense accumulates interest from the date it was incurred through to when the income associated with that cost is received. This ensures that expenditure is only incurred if the return exceeds the cost of interest.

Asset value is the value of all assets at the beginning of the interest period. The opportunity cost of investing in farm assets including livestock, machinery and infrastructure (sheds, yards etc) is captured in Farm Optimiser. Asset value operates in conjunction with interest rate to represent the opportunity cost of holding assets. Its role is to ensure that all assets that are selected have a return more than the interest cost, this ensures the optimal solution does not include assets that return less than investing the same money in a savings account (or to reduce core debt). This structure makes an equilibrium model generate a result similar to a multi-period model that accounts for the interest cost of money. Livestock flock structure is the main 'decision' that is altered by the inclusion of an asset value. For livestock, it ensures that the flock structure optimisation accounts for the opportunity cost of interest foregone from holding an animal till it is sold.

The interest rate for credit and debit are different for farmers' 'real money' in the bank. However, in Farm Optimiser the same interest rate is used to represent debit and credit. The reasons are:

1. Many farmers often have a core debt, so the farm cash position is usually negative even though their short term operating account may occasionally be positive. The differential interest rates are only justified if the farmer does not operate with a sweep facility to pay down core debt and then redraw when required later.

2. As discussed above, the asset value and the cashflow operate together in the optimisation of flock structure. This implies that the interest rate for the cash flow should be the same as the discount rate for the asset value.

Farm Optimiser tallies the total farm expenditure, adjusts it by a user defined return on expense factor and includes it in the objective to ensure the model achieves a minimum return on expenditure. The purpose of this is to represent farmer behaviour. It can also be used in the static equilibrium version to 'fudge' the risk associated with seasonal variation and reduce the optimal stocking rate to better align with on-farm values. The minimum rate of return on expenditure

(MinROE) is specified by the user and can be turned off. The current rate in the static equilibrium model (25%) was calibrated by a comparison of the model output with on-farm benchmarking (e.g. Planfarm, 2022).

There is no representation of a starting cash balance. If it is included, the model just selects the highest amount because that earns the most interest. The model can overdraw the working account if additional cash is required, so this does not affect the model solution.

Tax is also not represented for several reasons:

1. There are several mechanisms by which farmers seek to lessen their tax liabilities. Not all are 'economically rational' and not all are easily represented in a LP model.

2. Many farmers nowadays invest in farm management deposits and income-averaging as a means of taxable income averaging and to smooth working capital borrowings. If the FMDs could be used 'perfectly', then each year would have the same taxable profit. Thus, the optimal farm management is unaffected by the inclusion of tax.

3. Farm Optimiser is a bioeconomic model with the aim of optimising farm management. It is not a finance model.

Labour

To capture the dynamics of labour, the year is broken into labour periods (Rose, 2011). The supply of labour in each period by each labour source is calculated, and the labour required by each farm activity is determined and assigned to the given period/s.

In addition to the labour requirement described in the other sections of the model, there is a fixed labour requirement which reflects the labour required for administration tasks such as BAS, tax and pay roll, farm planning and upskill activities such as attending conferences or field days.

The amount of labour available in each period depends on the number of labour units and the hours worked each day. Labour can be supplied by three sources:

1. Casual staff – In the unrestricted model, casual staff can come and go at any time throughout the year as required. However, the user can fix the number of casual staff employed during each period of the year.

2. Permanent staff – Permanent staff work on the property all year (with an allocation for leave).

3. Manager staff (commonly the farm owner) – The farm manager works on the property all year. They control the overall farm plan and thus spend a fixed amount of time each quarter on farm planning, learning, record-keeping, purchasing and selling, and other office work.

Farm labour tasks can be allocated to a specific labour source where required. For example, farm planning must be completed by manager staff. Any labour source can complete unallocated tasks. To realistically reflect the labour hierarchy, casual and permanent staff both require a certain amount of supervision from the farm manager. The proportion of supervision is specified separately for seeding and harvesting. This is because during seeding and harvest it is likely that less supervision is required. Casual staff are generally less experienced and/or acquainted with the farm operation than permanent staff and thus require more supervision.

The importance of timeliness and the high labour requirement of seeding and harvest means staff often work longer days during those periods (Rose, 2011). To accommodate this, the user specifies the hours worked by each type of staff on the weekdays and weekends for both standard periods and seeding and harvest periods.

The farm manager and permanent staff have four weeks of holiday each year. The holiday timing is flexible (optimised by Farm Optimiser). This is because managers and permanent staff tend to have a less defined schedule, often taking multiple smaller holidays during the year or returning to the farm during holidays to check on things. Additionally, in Farm Optimiser, permanent and casual staff require supervision from the manager which means if the manager is forced to take their holidays in one big chunk the model may not be able to access labour resulting in inconsistencies if the period dates change. All labour sources take days off for Christmas, New Year's Day, and Easter. Permanent staff are also allocated a certain number of sick days per year. The user has the ability to alter the length and timing of worker leave.

Casual staff are paid on a per hour basis and the manager and permanent staff are paid an annual wage. All labour costs include superannuation and workers' compensation insurance.

Machinery

There is great variation in the type, age and investment in machinery between farms (Kingwell and Pannell, 1987). To account for this, the model accommodates a range of machinery options. The model user can then select which machinery complement is appropriate for their analysis. The machinery option selected determines the fixed cost, variable cost and machinery work rate. For the seeding activity, the land management unit can also impact the variable cost rate and work rate. For example, work rates are slower on heavy clay soils.

A machinery cost applies to all farm activities that require machinery. However, for both seeding and harvest, the work rate of the machinery affects the timelines of completion. For example, with smaller machinery, seeding takes longer, potentially incurring a late seeding yield penalty. Additionally, the model can hire contract services for seeding and harvest, although this can be limited by the user.

There are operating costs and depreciation costs associated with machinery. Operating costs refer to expenses incurred during usage such as for fuel, oil, grease, repairs and maintenance. Depreciation costs represent the decline in the value of the asset. Depreciation is made up of two components. Firstly, a fixed component which represents depreciation even if the machine is not used and secondly, a variable component which represents that asset value reduces faster with increased usage.

Uncertainty and short-term tactics

The two main sources of uncertainty in Australian farming systems are the high variance of world prices for most agricultural commodities (e.g. Hazell et al., 1990) and climate variability, which results in significant production variability (Feng et al., 2022, Laurie et al., 2019). To deal with short-term variability within the system, farmers implement tactical adjustments that deviate from the long-term strategic plan. Tactical adjustments are applied in response to unfolding opportunities or threats and aim to generate additional income or to avoid losses (Pannell et al., 2000). In Farm Optimiser price and weather variation is represented as a number of discrete options along with a range of relevant tactical adjustment options.

Weather variation

Weather uncertainty in Farm Optimiser can be included or excluded and the representation of uncertainty can be more or less detailed. Variability or uncertainty is represented using the modelling approach of discrete stochastic programming (Cocks, 1968, Rae, 1971, Crean et al., 2013). Discrete stochastic programming is a formulation of a decision tree. It requires the explicit specification of management choices and their possible consequences. The nodes or event forks are usually represented by a relatively small number of discrete outcomes. The inclusion of uncertainty allows management decisions to be made as the year unfolds (Norton and Hazell, 1986, Hardaker et al., 1991), which has been noted as an important aspect of farm management (Pannell et al., 2000, McCown et al., 2006). The three different Farm Optimiser frameworks are:

(i) A deterministic steady state expected weather-year framework (DSSE) (e.g. Kingwell and Pannell, 1987). In this framework the farming system is represented as a single discrete state that is statistically the expected weather-year. Representing a farm system by such a single state of nature requires use of expected inputs and outputs (e.g. the wheat yield is the average of all years). It assumes every year is the same and the finishing state equals the starting state. Thus, only strategic (or year-in year-out) management is represented and management does not change between years because there is only one branch of the decision tree being represented. This model includes 83,271 variables and 49,364 constraints.

(ii) A four-stage single-sequence stochastic programming with recourse (4-SPR) (e.g. Kingwell et al., 1991). A 4-SPR model represents the farming system as subject to a portfolio of discrete states of nature where each state represents a different type of weather-year that has separate or unique inputs and outputs to reflect different prices, weather conditions and production outcomes. All states begin from a common point that is determined by the weighted average of the end of all the weather-years, but then these states separate at various nodes during the production year to unveil the particular nature of that weather-year (Note: To minimise misrepresentation associated with the starting weighted average, the start of the weather-years is defined as the earliest season break). Once a weather-year has been identified, subsequent decisions are differentiated based on the known information about that given weather-year. For example, one node is the start of the growing season or 'break of season'. If that start is what is known colloquially as an 'early break', then after that starting point those types of weather-years can be managed differently to weather-years where the break occurs later. For example, in an early break it may be optimal to crop more area and run a higher stocking rate and vice-versa for a late break, although these decisions can only be made after the break of season is known. However, at the break of the season the subsequent conditions are uncertain (e.g. 30% chance of a poor spring and a 70% chance of a good spring). Thus, the decisions made at the break of season must factor in future uncertainty about the spring conditions. The 4-SPR model examines each possible outcome and its probability to determine the optimal decisions. These decisions are a suite of tactical adjustments made at each node that complement or adjust an overarching farm management strategy. The 4-SPR model is much greater in size, comprising 476,113 variables and 237,956 constraints.

(iii) A eight-stage multi-sequence stochastic programming with recourse framework (known as 8-SPR) (Xie and Huang, 2018). 8-SPR is similar to 4-SPR with the difference being that the discrete states represent a sequence of weather-years in equilibrium rather than a single year in equilibrium. Optimisation of management within the sequence of weather-years fully accounts for the temporal effects of management change between years. In Farm Optimiser, the production data in the 8-SPR is the same as the 4-SPR for the individual weather-years. The difference is that the 8-SPR framework more accurately represents carryover management implications from the previous year. For example, if stock were sold in the previous year the current year would start with a destocked position. This version of the Farm Optimiser model includes 4,571,881 variables and 2,140,700 constraints.

Price variation

There are two main methods to include price variation in whole farm LP.

1. Expected price variation (e.g. Kingwell, 1994): Expected price variation represents price variation by applying a discrete distribution to cashflow items after management decisions have been made. This method of representing price variation assumes that there is no knowledge of the price state, prior to purchasing or selling a commodity. The only known information is the expected price (i.e. a farmer does not know if they are in a high or low price year until they purchase or sell). Therefore, price variation has no impact on farm management for a risk neutral farmer. However, for a risk averse farmer price variation can alter their management. For example, if the grain price is more variable than livestock prices, it may be optimal for a risk averse farmer to have a higher livestock focus because it will reduce the variation in farm profit between years.
2. Forecasted price variation (Apland and Hauer, 1993): Forecasted price variation is a more realistic method achieved by including discrete states based on forecast information, allowing decision-making to change based on the forecasted conditions. The forecasted states are adjusted using a discrete distribution to reflect the actual prices received at purchase or sale. This requires a stochastic programming approach that increases model size and complexity.

Farm Optimiser currently uses method 1 because price variation has not been the focus of this doctoral thesis, Nonetheless, a likely valuable future improvement for Farm Optimiser would be to include forecasted price variation. Farm Optimiser’s flexible structure would facilitate inclusion of such price variation.

Currently, price variation is approximated in Farm Optimiser using a range of discrete price states for meat, wool and grain. The need to form discrete approximations of a continuous distributions is a necessary requirement for developing a LP model of farm management responses to price and weather-year states. By their nature, discrete stochastic programming models cannot consider all possible price states as described by continuous distributions. Rather continuous variables such as price need to be approximated by discrete states.

The price state scalars and their probabilities are calculated by fitting a multivariate normal distribution to historical prices, then summarises as discrete states by dividing the multi-dimensional probability density distribution into segments. A multivariate distribution is used so that correlations between commodities are accurately represented in the resulting price states. Grain and wool prices are better represented by log-normal distributions (Kingwell, 1996). Thus, before fitting the distribution, grain and wool data were subject to a log transformation. Additionally, the historical prices were CPI adjusted and detrended using a long-term moving average. The reason for detrending the price data was that the price states represented in Farm Optimiser serve the purpose of capturing yearly price variation (i.e. variations around the expected price for that year) rather than capturing within year price cycles.

Nonetheless, within year price cycles are accounted for in Farm Optimiser for products such as sale sheep that can be sold at different times during the year. Including the within year price cycles ensures that optimisation of the nutrition of sale sheep represents that sale data has an effect on expected price. Representing the annual price cycle also ensures that strategic management such as time of lambing is also evaluated correctly given the impact of time of lambing on likely turn-off dates.

To reduce model size and simplify input calibration, all meat classes (lamb, shipper, mutton, etc) receive the same meat price scalar. The same thing happens for classes of wool and types of grain. This simplification should not compromise the accuracy of the results because subclasses of a given commodity tend to have a high correlation (e.g. between 2000 and 2021 the correlation between light lamb and mutton was 96%). A further simplification was excluding price variation for input costs because input costs tend to vary less (Kingwell, 1996) and therefore the additional model size was not justified. The resulting assumptions are that all animal classes are 100% correlated, all wool microns are 100% correlated, all grains are 100% correlated and all input commodities have no variation. This assumption is not entirely accurate (e.g. canola and wheat prices are not 100% correlated) however, if in future analysis, price variation is of high importance this can easily be rectified by expanding the inputs.

Tactical management options

There are many tactical or adjustment options represented in Farm Optimiser that reflect a farmer's reality. The tactics are similar to, but an expansion of those represented by Kingwell et al. (1992) and revolve around land use area adjustment, land use inputs, whether a crop is harvested, baled or grazed as a standing crop, intensity of machinery use, labour utilisation, seasonal sheep liveweight patterns, tactical sale of sheep, grazing management of pasture and stubble, and supplementary feeding. The same tactical adjustments are made to all weather-years that are indistinguishable from one another at the time a tactical decision is implemented. Such weather-years are clustered at that decision point, as the node that later differentiates these weather-years is still in the future. By illustration, tactical adjustments selected at the early season break node have to be the same for all weather-years that have an early break, because at the time of making the break of season tactical decision the occurrence of follow-up rain and the spring conditions are unknown. Typical tactical adjustments include:

- Rotation phase - The area of each land use can be adjusted depending on the date of season break or other early indicators such as residual soil moisture from summer rainfall. Choice of rotation phase can also be delayed at the break of season, for example waiting to ensure it is not a false break. During this period of delay, pasture will germinate on these paddocks and is able to be grazed (the level of germination is dependent on the rotation history of the paddock). The potential for tactical adjustment of rotation phases depends on the land use history on each LMU because the choice for current land use is constrained by the land use history. Likewise, tactical adjustment affects subsequent rotation phase choice through its impact on altering the land use history provided.

- Land use inputs – In favourable weather-years additional chemicals and fertiliser can be applied to maximise yields and vice versa in poor weather-years. Note, in this analysis the input level for each land use on each land management unit in each weather-year is optimised by the user externally to the model, reliant on expert agronomist advice for the study region. The optimisation accounted for the clustering of the weather-years.

- Fodder crops - In adverse weather-years where either livestock feed is limiting or crops are frosted or are not worth harvesting, saleable crops can be turned into standing fodder. That is, instead of harvesting a crop it is grazed by livestock as summer feed.

- Bale crops - Crops planted with the expectation of being harvested for grain can be baled as hay. This may occur in adverse weather-years where either livestock feed is limiting or crops are frosted or are not worth harvesting.

- Labour supply - Permanent and manager labour is fixed (I.e. must be the same for all weather-years). However, casual labour can be optimised for each weather-year as it unfolds.

- Machinery contracting - If the timeliness of an activity is an issue, contract services can be selected to improve the work rate. This could be valuable in a late break weather-year to ensure the crops get the maximum possible growing season. Note, the assumption that contracting services are available can be changed.
- Dry seeding - A useful tactic to improve timeliness of seeding is to sow into dry soil, before the opening rains, to ensure crops experience the maximum possible growing season. If dry seeding is selected it is implemented for all weather-years that have yet to have the season break.
- Confinement feeding - Confinement feeding can be a good tactic to allow pasture deferment at the beginning of a growing season or to keep ground cover on paddocks in the late summer and autumn.
- Supplement feeding – In-paddock supplement feeding can be used as a tactic to help finish lambs for sale, ensure ewes reach target conditions for reproduction or to help meet energy requirements during weather-years with poor pasture growth.
- Changing liveweight - Altering livestock liveweight targets can be used as a tactic to handle varying feed availability due to seasonal variation e.g. animals can lose weight in poor feed years but this is associated with lower production per head.
- Not mating ewes - If the feed supply is sufficiently poor prior to joining then there is the option of not mating ewes. This might be most relevant if mating ewe lambs.
- Selling scanned dry ewes or other ewes at scanning – Sale of dry sheep can be a useful tactic if the year is unfolding unfavourably.
- Retain dry ewes - If the strategy is to sell dry ewes, and the weather-year is favourable, a tactical adjustment can be to retain the dry ewes until shearing, thereby generating wool income and then a further decision is to retain them for mating the following year.
- Selling at other times – The ewes and lambs' sale time can be adjusted with the value received depending on the liveweight and condition of the animals at sale. In Farm Optimiser there are ten selling opportunities throughout the year for ewes and eight sale opportunities for lambs and wethers.

Objective

The objective of the model is to maximise expected utility. In the case of risk neutrality, expected utility is equivalent to expected profit, meaning that the optimal farm management plan is that which maximises farm profit. In the case of risk aversion, utility increases at a diminishing rate as profit increases. Thus, when farm profit is low, an extra dollar of profit provides more utility than when farm profit is high. This means a risk averse farmer aims to reduce profit variation (i.e. increase profit in poor years at the cost of reduced profit in the good years). For example, if the crop and stock enterprise on the modelled farm are similar but grain prices are more volatile, then risk aversion will shift resources towards the stock enterprise to reduce risk (profit variation).

Constant absolute risk-aversion (CARA) and constant relative risk-aversion (CRRA) are two well known utility functions. Both have been previously used in stochastic farm modelling (Kingwell, 1994, Kingwell, 1996). Both methods are included in Farm Optimiser. CARA is a negative exponential curve: $U = 1 - exp(-a * x)$ where U is utility,a is the Pratt-Arrow coefficient of absolute risk aversion and x is the return to management and capital. The Pratt-Arrow coefficient is a user input that controls the level of risk aversion. Kingwell (1994) used two levels: 0.000 003 and 0.000 005 to represent moderate and high levels risk-aversion. CRRA is a power function denoted by: $U = \frac{W^{(1-R)}}{1-R}$ where U is utility, W is terminal wealth and R is the relative risk aversion coefficient. The relative risk aversion coefficient is a user defined input that controls the level of risk aversion. Kingwell (1996) used values within the range of 0.1 to 3.0 to represent low to high levels of risk-aversion.

Both methods have limitations, most of which can be minimised if the modeler is aware. A CARA specification implies there are no wealth effects on a farmer's income and price security decisions. In practice, the CARA specification means that the farmer's risk management decisions, particularly in favourable states of nature (e.g. good weather-years with high commodity prices) when a farmer's wealth is boosted, will be different and more concerned with income stability than those that would arise with a CRRA specification. The limitation of the CRRA method is that it cannot handle a negative terminal state. Additionally, because CRRA is impacted by terminal wealth, MINROE and asset opportunity cost (discussed in the finance section) will affect the impact of risk aversion, which is not technically correct because these are not real costs incurred by the farmer.

The utility functions discussed above are non-linear. To accommodate this in Farm Optimiser, a piecewise technique is applied which approximates the function using 13 linear segments.

Validation

As an optimisation model, the sort of validation strategies used for simulation models are not applicable. However, significant time has been spent by the authors with local farming experts examining the outputs of each module to ensure that results and behaviour of the model are realistic and well aligned with actual farms in the region. Additionally, Farm Optimiser was constrained and benchmarked against regional data to ensure that for a given farm structure the results are realistic. Furthermore, Farm Optimiser builds on many the of fundamentals used in MIDAS which have been extensively tested in Western Australia for around 40 years since its creation (Kingwell and Pannell, 1987).

Table 1: Comparison of regional benchmarking results with Farm Optimiser output when stocking rate, pasture area and prices in Farm Optimiser were fixed to the benchmarking averages.

	Sheep gross margin[1] ($/WgHa)	Average wool cut (kg/hd)	Weaning %	Cropping gross margin [1]($/Ha)	Barley Yield (t)	Canola Yield (t)
Regional[2]	$652	4.03	99	$570	4.2	2
Farm Optimiser	$620	4.15	100.5	$549	4	1.8

[1] Only includes variable costs.

[2] Regional results were from 2021, an above average production year (livestock prices were very favourable in 2021)

Conclusion

Farm Optimiser is a detailed whole farm mathematical program consisting of numerous modules that represent each aspect of the farm system. The level of complexity included in each module is part of the art of building whole farm model. Enough detail must be included to ensure the research topic being evaluated is fully captured. However, it is often also subject to resource allocation and parameterisation data. For example, in the last decade, Farm Optimiser and MIDAS have been more widely used in the livestock industry and as such the livestock and pasture representation in the model is more detailed than cropping. In the future, modules can easily be improved when the resources and data are available. For example, nutrient response curves could be added to the cropping modules that allow fertiliser application to be optimised or cattle could be included.

Overall, the development of Farm Optimiser has been a big step forward for the farming industry. Farm Optimiser is the first Western Australian whole farm optimisation model built in Python. Utilising modern programming concepts, it is able to represent a high level of farm biological, technical and economic details, including weather and price variation. This opens the doors to many different applications by model users. Farm Optimiser has the capacity to answer questions that were not previously possible such as what is the importance of single year and multi-year stochastic farm modelling versus steady state farm modelling. Furthermore, Farm Optimiser has several sets of inputs allowing it to be applied in different regions. Taking Farm Optimiser to a new region requires the user to calibrate the relevant model inputs which can be a significant job (e.g. Young et al., 2020). However, with the use of simulation modelling, expert advice, research data, benchmarking data, on-farm data and satellite data this is an achievable task.

Additional features of Farm Optimiser include segregated inputs, cloud integrated, version control and automated documentation. All these functions make Farm Optimiser highly flexible, more scalable, more transparent and more user friendly.