

Choosing crop rotations under uncertainty: a multi-period dynamic portfolio optimization approach

Geoffrey Lee^a, Chenming Bao^a, Nicolas Langrene^a and Zili Zhu^a

^a *Digital Productivity Flagship, CSIRO*
Email: geoffrey.lee@csiro.au

Abstract: To achieve good crop yields, farmers are aware of the importance of good rainfall during the growing season. However, their choice of crop to plant may not necessarily be the optimal choice when climate uncertainty exists. Indeed, current cropping allocations may be driven more by historical trends and tradition than by the future forecasts of climate scenarios and commodity prices. In a digital agricultural future, farmers may instead use information on crop yields and water usage to make cropping decisions each year that optimize their cash flow under the uncertainty of climate conditions and commodity prices.

In this paper, we apply dynamic portfolio optimization techniques to the development of a simulation-based numerical method for making dynamic optimal cropping decisions. This method relies on a backwards recursive approach developed to solve the American option pricing problem. At each time step backwards from the end of the decision time period, the optimal expected future cash flow, or the so-called continuation function, is determined by using the Least Squares Monte Carlo method. As an example, we use a representative farm in Australia with four paddocks that can grow wheat, rice, barley and canola, and we also regard the corresponding commodity prices as stochastic variables. We compute the optimal crop rotations each year under different rainfall scenarios that maximise the expected utility over a fixed time period of 20 years. We evaluate the performance of the dynamic cropping strategies by comparing the expected value and standard deviation of future cash flows against those generated from static cropping strategies.

Keywords: *Crop rotation, rainfall, portfolio optimization, Least-squares Monte Carlo, approximate stochastic dynamic programming*

1 INTRODUCTION

Historically, farmers have chosen crop rotations on individual farms through cultural knowledge and tradition, and, in some cases, also through national regulations. However, the uncertainties inherent in farming, such as future commodity prices and the increasing threat of climate uncertainty may render these choices ineffectual. For example, decreasing commodity prices may cause a particular crop to cease to be financially viable. Furthermore, increased climate variability may render other crops impossible to grow, due to prolonged drought or extreme weather events such as cyclones and floods. Under these scenarios, entire regions could experience a shortage of certain crops, leading to a risk of damaging food security in the region (Kandulu et al. (2012); Huh and Lall (2013)).

To mitigate the risk caused by uncertain commodity prices and climate conditions, farmers should introduce diversity of crop production to ensure future profitability. For example, a study by Kandulu et al. (2012) found that introducing diversity of crops under climate uncertainty reduced the standard deviation of returns and increased the probability of breaking even. However, this study only assumes that crops are equally allocated, and does not allow for unequal allocations of crops. While this diversification may improve the risk, it may not be the allocation that optimizes the returns, for example, over a period of time.

Several studies have examined the optimal cropping allocation problem at a large scale, for example Kumar and Pathak (1989); Havlik et al. (2010); Fuss et al. (2011). The crop allocation problem can also be formulated as a stochastic programming problem, which can be solved using methodology developed for portfolio optimization in the finance industry. In particular, the farmer can be seen as investor, wanting to maximise their returns (or other objective function) by allocating their land to growing different crops. These crops have stochastic commodity prices, and so the farmer can be seen to be investing in risky assets in the face of stochastic risk factors. A study by Huh and Lall (2013) treats the crop allocation problem in this manner, examining the allocation of land between two crops with differing water requirements under the uncertainty of future rainfall and commodity prices.

In this paper, we treat the crop allocation problem as a stochastic *dynamic* programming problem, and adapt the Least-Squares Monte Carlo (LSM) method (Longstaff and Schwartz (2001); Stentoft (2004)) initially developed for valuing American-style options to solve this problem. This paper will use the LSM algorithm to determine the optimal cropping rotations of a typical farm over a period of 20 years under the uncertainty of future commodity prices and for three different rainfall scenarios. The rest of the paper is organized as follows. Section 2 sets up the optimization problem and describes the LSM methodology we employ to solve it. Section 3 details the case study of a single farm we use to illustrate our optimization technique. Finally, Section 4 presents our numerical results, while a summary is provided in Section 5.

2 CROP ALLOCATION AS AN OPTIMIZATION PROBLEM

Part of the work a farmer must perform is deciding which crops to plant and where (if the farmer has multiple paddocks) to ensure that a good yield is obtained and thus sufficient income is made. If the farmer has a choice of multiple crops to allocate to their land, the allocation problem can be seen as analogous to a portfolio allocation problem. We therefore treat the farmer as an investor, with N crops, or assets to allocate to their land. The farmer's decision is going to be influenced by the money made from growing crops, which is clearly dependent on the commodity price. We denote the commodity price of the i -th asset at time t by S_t^i , and we model each commodity price as a stochastic process in the probability space $(\Omega, \mathcal{F}_t, \mathbb{P})$.

We then assume that the farmer would like to optimize the cropping allocation to, for example, maximise the income generated by selling the crop yield over a time period of T years, also assuming that the farmer can perform a crop rotation annually. This leads to the need for a dynamic solution to the portfolio optimization problem. We note that other objectives can be added to be optimized, such as minimising the water consumption, or maximizing the carbon sequestration.

We can therefore formulate a stochastic dynamic programming description of the problem as (Bao and Zhu (2013); Bao et al. (2014)):

$$V_t(\omega_t|x_{t-1}) = \max_{x_t} \left\{ \mathbb{E} [f_t(x_{t-1}, x_t, \omega_t) + V_{t+1}(\omega_{t+1}) | \mathcal{F}_t] \right\}, \quad (1)$$

$$\omega_t \in \mathcal{F}_t$$

where x_t is the vector of portfolio weights (i.e., the weighting on each crop) for the realized commodity price

vector ω_t at time t and V_t is the value function at time t . The value function can be seen as the expected total future utility at time t given that all future portfolio weights are chosen optimally.

The function $f_t(\cdot)$, termed the utility function, represents the risk preference of the farmer. Instead of simply choosing the crop with the highest price to maximize the income, farmers may want to take less risk by choosing crops that have a higher probability of giving some income, if not the greatest income. In this work, We consider different risk preferences of investors by using the Power utility function which is given by:

$$u(w) = \frac{w^{1-\alpha}}{1-\alpha}, \quad \alpha > 0. \tag{2}$$

The Power utility function is also known as the constant relative risk aversion (CRRA) utility function. It's a relative measure of risk aversion, defined by $-wu''(w)/u'(w)$, where α is a constant. A larger value of α leads to a more conservative outlook – that is, the investor values a smaller risk over higher returns.

In order to solve the stochastic dynamic programming problem characterised in Equation 1, we use the Least Squares Monte Carlo (LSM) method initially developed to solve the American option pricing problem. A detailed description of the LSM method can be found in the works of Bao and Zhu (2013) and Bao et al. (2014). Briefly, the LSM method employs a backwards recursion technique, solving the problem from the last time step to the first. The LSM method proceeds by approximating the conditional value function given in Equation 1 using a least-squares regression onto appropriate basis functions. For this work, we use polynomial basis, including terms to third order.

3 CASE STUDY

For a comprehensive investigation into the application of stochastic dynamic portfolio optimization methodology to the agricultural space, a detailed study of the convergence properties of the algorithm, including comparisons to appropriate benchmarks, is required. However, this treatment is beyond the scope of this work, and will be left to future studies. In this work, we seek to illustrate the advantages of using stochastic dynamic portfolio optimization by means of a simple case study. We consider an example of a typical farm with four paddocks available for growing crops. This farm can plant four different crops: barley, wheat (Australian standard white), rice and canola.

We assume that only one crop type can be grown in each paddock, implying that the fractional amount of land each crop can occupy is discretised and can only change by quarters leading to the possible fractional values of (0, 0.25, 0.5, 0.75, 1). Ensuring that crops cannot occupy more than the total land size results in 35 possible cropping allocations. The profit generated per unit of land from implementing these cropping allocations is found by multiplying the i -th commodity price S_t^i in dollars per tonne by the commodity's yield in tonnes per hectare.

3.1 Stochastic crop price models

Rather than modelling the commodity (crop) prices directly, we instead choose to model the log-returns R_t at time t . Denoting the i -th commodity price valued at time t as S_t^i , where i is a component of the current risk factor set: $\Psi = \{\text{barley, wheat, canola, rice}\}$, R_t is the four dimensional log-return rate of risk factors S_t^i , $i \in \Psi$:

$$R_t = \left[\ln \left(\frac{S_t^i}{S_{t-1}^i} \right) \right], \tag{3}$$

We assume a 4-dimensional mean-reverting model as the canonical model for R_t , such that

$$R_t = R_{t-1} + \Xi[R_{t-1} - \mu_R] + \Sigma'W_t, \tag{4}$$

where Ξ and Σ' are 4x4 real value matrices, μ_R is the long-term mean and W_t is a four-dimensional Brownian motion with zero mean vector and a 4x4 unit covariance matrix.

In this study, the parameters μ_R , Ξ and Σ of the canonical model in Equation (4) are estimated through a standard maximum likelihood algorithm on a historical market dataset. We use crop prices from the document Australia Commodity Statistics 2014 published by the Australian Bureau of Agricultural and Resource Economics and Sciences ABA (2014). More specifically, we use: (1) "Indexes of prices received by farmers in Australia" for wheat, barley and canola (pp.19); (2) "Gross unit value of production" for rice price (pp. 183).

The full range of data from ABA (2014) runs from the 1981/1982 financial year through to the 2012/2013 financial year representing 32 observations for each crop.

We note that we also assume a constant bond yield, and use the Australia 10 Year bond yield which is 2.36% at the time of this study (April 2015).

3.2 Yield model

Following the work of French and Schultz (1984) and Sadras et al. (2003), we assume that a crop's yield is linearly proportional to the growing season rainfall (GSR), where the constant of proportionality is termed the Water Use Efficiency (WUE). In particular, we take the yield to be

$$Y = (GSR - E + ASW)WUE, \tag{5}$$

where evaporation E is the amount of water that is lost from the soil surface over the growing season due to surface evaporation and Available Soil Water ASW is the amount of plant available water present in the crop rooting zone of the soil when the crop was sown. The WUE is the amount of grain that a particular crop has the potential to produce for each mm of available water per hectare of land. A range of values for the WUE of various crops exists in the literature Sadras et al. (2003); Singh et al. (2005); Sadras and McDonald (2012), so we use representative values presented in Table 1.

Table 1. Crop water use efficiency

| Crop | Rice | Wheat | Canola | barley |
|-------------------------------|------|-------|--------|--------|
| Water use efficiency kg/mm/ha | 7.1 | 14.3 | 9.5 | 14.5 |

We approximated the GSR by annual rainfall, as multiplying the annual rainfall by a constant between 0 and 1 to obtain the GSR would result in a constant shift in yields. We obtained the past x years' annual rainfall from the Bureau of Meteorology website to obtain a data sample for the GSR. We then fitted the GSR data sample to an ARIMA time series family. Based on the Akaike information criterion (AIC) and the Bayesian information criterion (BIC), a null model (ARIMA(0,0,0) with non-zero mean) is the best fit model. The hypothesis that the GSR data sample are white noises cannot be rejected.

We use a linear model as the forecast model:

$$GSR_{t+1} = (1 + (t + 1)\gamma^{\{i\}})\mu_{GSR} + \sigma e_{t+1}, \tag{6}$$

where μ_{GSR} is the long term mean value of the GSR, e_{t+1} is a white noise and $\gamma^{\{i\}}$ is a hypothesis trend for each of the forecast rainfall regimes. We consider a rainfall scenario with $\gamma = 2\%$, i.e., that the mean GSR decreases 2% per annum.

Finally, we penalise the use of a given paddock by the same crop in consecutive years to simulate the need to fertilise the paddock. Due to computational considerations, for this study we choose a single cost of \$20,000 for growing the same crop in a given paddock in consecutive years. Future work will investigate the effect of changing this penalty on the optimal decisions and expected future cash flows.

4 RESULTS AND DISCUSSIONS

Due to computational considerations, we use a total of 10000 Monte Carlo sample scenarios of the risk factors (crop prices and rainfall) to compute the expected value functions for this case study. In future studies, we will investigate the effect of larger sample numbers. We split the 10000 sample scenarios into a 5000 sample set used as a training set to compute the Least-Squares coefficients used in constructing the continuation value, and the remaining 5000 sample scenarios are used as an out-of-sample test to determine the optimal cropping strategies.

We then compare the overall performance of the dynamically chosen crop rotations to statically chosen cropping allocations by looking at the expected value and standard deviations of the future cash flow. Table 2

Table 2. Summary of cash flows for CRRA- α dynamic crop rotations and static weighted crop rotation.

| Type | Rotation | Mean (m\$) | SD (m\$) |
|---------|-------------|------------|----------|
| Static | 100% Rice | 4.7193 | 3.1403 |
| | 100% Wheat | 5.9064 | 3.9443 |
| | 100% Barley | 7.6286 | 6.4953 |
| | 100% Canola | 3.7165 | 2.0753 |
| | 25% Each | 5.7696 | 3.2883 |
| Dynamic | CRRA3 | 8.1906 | 5.6614 |
| | CRRA5 | 7.6618 | 4.8585 |
| | CRRA7 | 7.5156 | 4.6434 |
| | CRRA9 | 7.4420 | 4.5091 |
| | CRRA11 | 7.3980 | 4.4347 |

presents both the expected future cash flows and the standard deviations for a selection of dynamic and static cropping rotation schemes, while Figure 1 presents a mean-sigma plot, where the expected future cash flow is plotted against the standard deviation, for both the static and dynamic cases. All static cropping rotations are plotted, as well as the dynamic results for CRRA α values from 3 to 15.



Figure 1. Interpolated efficient frontier for dynamic and static rotations on mean-sigma plot.

Examining the static allocation results presented in Table 2, creating a portfolio of crops by allocating 25% of the land to each type of crop does indeed reduce risk; while the expected future cash flow is less than cropping 100% wheat or barley, the standard deviation of the diversified allocation scheme is less than either of wheat and barley.

Moving to the dynamic allocation results, we can see that as the CRRA parameter α is increased, the expected

future cash flow reduces, but so does the standard deviation. This demonstrates that the parameter α does indeed indicate the risk preference of the investor. However, with the CRRA utility function, a limit on the standard deviation is reached. As α is increased, the standard deviation converges towards a value of 4.5. A different choice of utility function may enable a greater range of standard deviations to be achieved. This is a topic for future research.

We also observe from Figure 1 that the blue curve corresponding to the dynamic allocation results appears to be shifted upwards from the red curve indicating the static results. This shift results in two related benefits to farmers using the dynamic cropping allocation scheme. Firstly, the farmer can receive the same expected future cash flow with less risk than when using a static allocation scheme. This corresponds to drawing a horizontal line through Figure 1. More specifically, in Table 2, we observe that the static allocation scheme that gives the largest expected future cash flow is planting 100% of the land with barley, with an expected future cash flow of \$7.6286 million, and a standard deviation of \$6.4953 million. We compare this to the dynamic strategy CRRA 5, where the expected future cash flow is similar at \$7.6618 million, but the standard deviation is only \$4.8585 million, a difference of \$1.6368 million.

Secondly, for a given level of risk (i.e. standard deviation of future cash flow), the farmer can expect to yield greater future cash flows using dynamic allocations when compared with static allocations. This corresponds to drawing a vertical line in Figure 1.

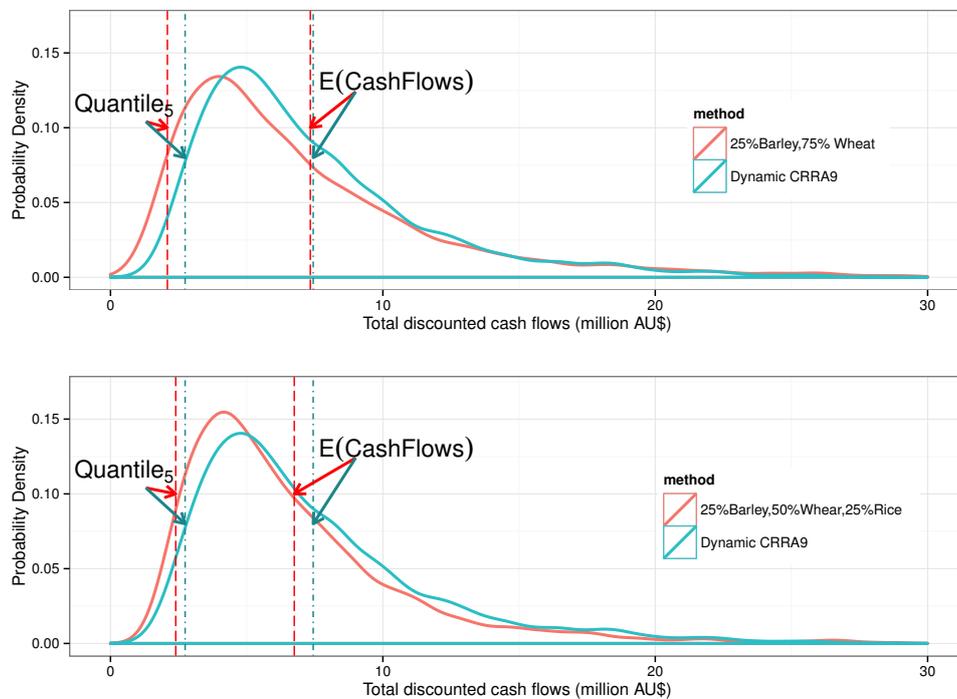


Figure 2. Probability density functions for two static cropping strategies and the CRRA 9 dynamic cropping allocation strategy.

The two related benefits farmers can gain in using dynamic cropping allocations are reinforced by looking at the probability distributions of future cash flows, as seen in Figure 2. The top panel compares the static cropping strategy of allocating 25% of the land to barley and 75% of the land to wheat to the dynamic strategy using the utility function CRRA 9. Each of these strategies give very similar expected cash flows (the vertical lines markeg $E(\text{Cash Flows})$ in Figure 2), but the static strategy has a wider distribution, as indicated by the position of the 5% quartiles for each strategy. The bottom panel displays a comparison between the same dynamic strategy using the CRRA 9 utility function and the static allocation of 25% barley, 50% wheat and 25% rice. Here, the 5% quartiles indicating the spread of the distribution are similar, but the dynamic cropping strategy gives a higher return.

Furthermore, Figure 2 indicates a further benefit of the dynamic cropping strategy over static strategies. In both panels, the dynamic strategy's probability distribution indicates that there is less chance of receiving small expected cash flows; that is, the dynamic cropping strategy limits the downside risk in comparison to static strategies.

5 CONCLUSION

We have demonstrated the application of stochastic dynamic programming to the problem of allocating crops to maximise the investors' expected utility under the uncertainty of future crop prices and climate conditions. We utilised the Least-Squares Monte Carlo method to solve the stochastic dynamic programming problem. We found that the dynamic optimization method can achieve an expected future cash flow \$0.8077 million more than the best static allocation result. We also found that a farmer allocating their crops dynamically can either obtain a greater future cash flow for the same level of risk, or alternatively, achieve the same expected future cash flow with less risk. In addition, the farmer can also reduce the downside risk of small expected cash flows by allocating their crops dynamically.

REFERENCES

- (2014, December). *Agricultural commodity statistics 2014*, ABARES. Australian Bureau of Agricultural and Resource Economics and Sciences.
- Bao, C., G. Lee, and Z. Zhu (2014). A simulation-based portfolio optimization approach with least squares learning. In *Proceedings of the World Congress on Engineering*, Volume 2.
- Bao, C. and Z. Zhu (2013). Land use decisions under uncertainty: optimal strategies to switch between agriculture and afforestation. In *MODSIM2013*, pp. 1419–1425. Modelling and Simulation Society of Australia and New Zealand.
- French, R. and J. Schultz (1984). Water use efficiency of wheat in a mediterranean type environment ii: Some limitations to efficiency. *Australian Journal of Agricultural Research* 35(765-775).
- Fuss, S., P. Havlik, J. Szolgayova, E. Schmid, and M. Obersteiner (2011). Large-scale modelling of global food security and adaptation under crop yield uncertainty. In *EAAE 2011 International Congress*. European Association of Agricultural Economists.
- Havlik, P., U. Schneider, E. Schmid, H. Bottcher, S. Fritz, R. Skalsky, K. Aoki, S. D. Cara, G. Kindermann, F. Kraxner, S. Leduc, I. McCallum, A. Mosnier, T. Sauer, and M. Obersteiner (2010). Global land-use implications of first and second generation biofuel targets. *Energy Policy*.
- Huh, W. and U. Lall (2013). Optimal crop choice, irrigation allocation, and the impact of contract farming. *Production and Operations Management*.
- Kandulu, J., B. Bryan, and J. Connor (2012, July). Mitigating economic risk from climate variability in rain-fed agriculture through enterprise mix diversification. *Ecological Economics* 79, 105–112.
- Kumar, R. and S. Pathak (1989). Optimal crop planning for a region in india by conjunctive use of surface and groundwater. *International Journal of Water Resources Development* 5(2), 99–105.
- Longstaff, F. and E. Schwartz (2001). Valuing american options by simulation: A simple least-squares approach. *Review of Financial Studies* 14(1), 113–147.
- Sadras, V. and G. McDonald (2012). Water use efficiency of grain crops in australia: principles, benchmarks and management. Technical report, Grains Research and Development Corporation.
- Sadras, V., D. Roget, and M. Krause (2003). Dynamic cropping strategies for risk management in dry-land farming systems. *Agricultural Systems* 76, 929–948.
- Singh, R., J. Mullen, and R. Jayasuriya (2005). Farming systems in the murrumbidgee irrigation area in nsw economic research report no. 10. Technical report, NSW Department of Primary Industries.
- Stentoft, L. (2004). Assessing the least squares monte-carlo approach. *Review of Derivatives Research*.