Modelling water blending - sensitivity of optimal policies

Webby, Brian and David Green and Andrew Metcalfe

School of Mathematical Sciences, Faculty of Engineering, Computer and Mathematical Sciences
The University of Adelaide, Adelaide, SA 5005, Australia.
E-Mail: roger.webby@adelaide.edu.au

Keywords: *Integer programming, linear programming, stochastic dynamic programming, CVaR*

ABSTRACT

Throughout the world there is increased demand on freshwater resources. Such resources are limited in quantity and erratic in availability but are renewed over time through the water cycle. Planning for water management then has temporal scales and stochastic variation to consider. There may also be several sources of water suitable for non-potable purposes. Supplying these by supplementing higher quality water with sources of differing quality and availability is the blending problem studied in this paper.

Consider a supplier in a non-potable use water market. The supplier obtains water from a range of sources and delivers it to a number of users. One such application, on which our model is loosely based, is the integrated water resource management project of the northern region of metropolitan Adelaide. Here the major potential sources are the stormwater harvesting and aquifer storage project of the City of Salisbury, recycled water from Bolivar water treatment plant and Adelaide's reticulated, potable water supply. Potential major users of the water (sinks) are a wool processing plant, a residential grey water network and council parks and gardens (Figure 1).

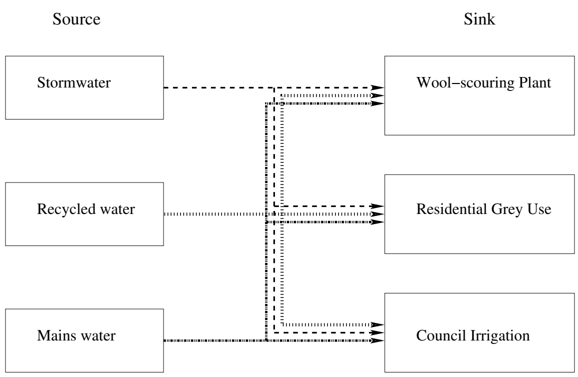

Figure 1. Sources and sinks for the water allocation problem

This is an optimisation problem of blending water from various sources to meet quantity and quality requirements of the sinks, with the objective of maximising expected profit from undertaking supply. The supplier undertakes to supply a guaranteed quantity of water to each sink in each time period and charges a premium for this water. There are further, preferred demands which are delivered at the supplier's discretion and for which a lesser cost applies, both amounts to meet salinity conditions. Profit is maximised if lowest cost sources are used to supply the firm demand and, perhaps, the preferred demand. The monetary value of the water resources is assessed by linear programming (LP) and by integer linear programming (ILP).

The problem is solved over sequential time periods, that is, it is solved as a stochastic dynamic programming (SDP) problem with stormwater as a stochastic source. The decision variable in this problem is the relative proportion of water from each source supplied to each user at each time step. Identifying the optimal decision for each state of the system gives a policy for the optimal use of available stormwater, given the probabilities of replenishment of this resource. SDP usually has the objective of maximising expected monetary value (EMV) but here we illustrate the use of an alternate objective, that of minimising exposure to risk of loss from undertaking supply, a risk quantified by the Conditional Value-at-Risk (CVaR). We compare solutions found by maximising EMV and by minimising CVaR.

The initial model broadly discretises the state space and we later assess solutions from using finer discretisations to represent two continuous variables - time and volume. The problem is partly one of unit commitment because stormwater and recycled water, lacking an extensive reticulation network, are likely to be handled in a water market in discrete, relatively large amounts. Then an integer program may be more realistic, but linear programs are easier to solve. We find and compare LP and ILP solutions for the problem.

We describe and comment on the policies that our model generates, giving details of the long-term results of particular strategies. The profits found by the LP and the ILP approaches differ slightly but draw closer with increasing discretisation of the state space.

1 INTRODUCTION

Water blending is likely to become much more common as demand for limited resources increases. Then techniques of mathematical analysis used for the similar problem of scheduling of reservoir releases, for example, Archibald *et al.* (2006) and Cabero *et al.* (2005), could be applied in a water market. We present a model for non-potable use, supplying water from 3 sources to meet demands at 3 sinks. We use LP and ILP to determine the optimal blend of water between source and sink, and couple these results to an SDP that finds the long-term optimal policy for use of water from a stochastic source.

The sources are storm, recycled and mains water, of which stormwater has the lowest salinity. Thus it is a preferred source for blending but is erratic in quantity and variable in availability with time. Recycled water is more constant in availability and quantity but has a high concentration of salts. Mains water is between the other two sources in salinity and has relatively large and constant availability, and is treated to the standards for potable use. Given the limited availability of water from other sources or its inability to meet quality requirements, mains water can be blended to meet demands for non-potable use.

Of the 3 sinks, the wool processing plant has a demand for good quality water; the grey water network takes water of a lower standard; the council can use water of a lower quality again. We focus on the salinity concentration of the various water sources here but other relevant water quality issues which could be included in the mathematical programming given suitable data include pH, bacterial levels and mineral concentration. The maximum acceptable salinity levels are typical for their intended use but also represent more generally water suitability for a particular range of uses. For example, water which has salinity levels acceptable to the wool processor can also meet drinking quality standards in that regard. Water suitable for non-potable home use would be generally acceptable for commercial cleaning and garden watering. Water suitable for council use, for example, irrigation of recreation areas, would be restricted to use in hardier environments.

The price of stormwater should reflect its cost of capture, cleaning and storage. However, these costs may be somewhat defrayed since the capture of stormwater runoff avoids a certain amount of damage to near-coastal environments. Similarly, a reduced price for recycled water may reflect the cost of its disposal to the environment, encouraging its reuse. Mains water is processed to a high standard and water from this source is priced accordingly.

One approach to discretising continuous variables is to identify typical states that a system has been observed in in the past, thereby enabling the system to be represented in low-dimensional terms (Archibald *et al.* (2006)). Our initial model demonstrates the specification of the LP/ILP - SDP model but a realistic application would use a larger state space. However, the conclusions from comparing the LP and ILP solutions are shown to carry over to an expanded state space.

A water utility may well be managed with a risk averse priority and we compare optimal policies found by a profit maximisation (EMV) objective and by CVaR. Conditional Value-at-Risk is a risk measure developed in the financial services arena. It is defined as the expected loss given that the loss is greater than or equal to a quantile of the loss distribution called Value-at-Risk (VaR) (Rockafellar and Uryasev (2002)). VaR is the maximum loss expected to be incurred over a given time horizon at a specified level of probability α. For cumulative probability distribution for loss, $F(d, z)$ where d is a decision, z is loss and E is expectation

$$\mathrm{CVaR}_{\alpha}(d) = E\{\, z \mid F(d, z) \geq \alpha\}. \qquad (1)$$

CVaR has been assessed for use in the electricity generation market (Cabero *et al.* (2005)) as a risk measure suitable for developing rules for the optimal allocation of stochastic resources.

2 MODEL SPECIFICATION

2.1 Water Characteristics

Water is available from three sources: stormwater from a reservoir which could be holding dams and/or aquifer storage sites (s); recycled water from a sewage treatment plant (r); and mains water (m). Let A_i represent the quantity of water available from each source i where $i \in \{s, r, m\}$. The available quantity of stormwater is treated as a stochastic variable, whereas available quantities of recycled and mains water are treated as deterministic (Table 1).

We consider volumes of water in terms of discrete units, and time (t) is modelled as discrete with $t \in \{0, \dots, T\}$. Our initial model assumes that the stormwater stored in the reservoir $A_s \in \{0, 1, 2\}$ and a time step of 1 year. During the year there may be 0, 1 or 2 units available from the reservoir if capacity is available. The amounts of rainfall/runoff are

- 0 -"low" - streamflows are limited and no water is taken from the environment
- 1 -"fair" - flows are moderate and 1 unit can be stored if capacity is available

- 2 -"wet" - flows are above average and 2 units can be stored if capacity is available.

There are 3 sinks: a wool processing plant (w), an urban grey water network (u), and the council (c). Each sink (j) has a firm demand M_j and a preferred demand $P_j > M_j$. Water supply above the firm demand is sold at a lower price because its supply is not guaranteed (Table 1). Let $x_{i,j}$ be the amount of water supplied from source i to sink j where $i \in \{s, r, m\}$ and $j \in \{w, u, c\}$.

We assume typical salinity levels for stormwater of 100 mg/l or ppm, for recycled water of 2000 mg/l and for mains water of 500 mg/l. We set the maximum acceptable salinity level for the wool processing plant at 500 mg/l, for urban grey water supply at 900 mg/l and for council purposes at 1300 mg/l. Let S_i be the salinity of each source and let B_j be the maximum tolerance for salinity at each sink.

Table 1. Water characteristics of sources and sinks. Availability, firm and preferred demand are in units of volume, salinity in mg/l. Stormwater availability is a random value $\in \{0, 1, 2\}$.

source:	**mains**	**recycled**	**storm**
availability	up to 20	up to 5	{0, 1, 2}
salinity	500	2000	100
sink:	**wool**	**urban**	**council**
firm demand	2	3	3
pref. demand	3	6	6
max. salinity	500	900	1300

Costs and returns per unit of water supplied are shown in Table 2. Let c_{ij} be the costs accrued from supplying 1 unit of water from source i to sink j. Let G be the return per unit for supplying the guaranteed amount and R be the return per unit for supplying the preferred demand.

Table 2. Costs (c_{ij}) and returns (G, R) per unit of water supplied ($)

	storm	**recycled**	**mains**
wool	1300	900	2550
urban	1300	900	2550
council	1300	900	2550
return for firm demand			4000
return for preferred demand			2500

2.2 The linear program

The LP finds the optimal allocations of water between the various sources and sinks with the objective of maximising profit while meeting availability and salinity constraints. The linear program is:

$$\max \left[G \sum_j M_j + \sum_{i,j} (x_{ij} - \sum_j M_j)\, R - \sum_{i,j} x_{ij} c_{ij} \right]$$

$$\begin{aligned} \text{such that} \quad \sum_i x_{ij} &\geq M_j && \text{for } j = w, u, c \\ \sum_i x_{ij} &\leq P_j && \text{for } j = w, u, c \\ \sum_j x_{ij} &\leq A_i && \text{for } i = s, r, m \\ \sum_i S_i\, x_{ij} / \sum_i x_{ij} &\leq S_j && \text{for } j = w, u, c \\ x_{ij} &\geq 0 \end{aligned}$$

The ILP is the same except for the last constraint which would read $x_{ij} \in \mathbb{N} \cup \{0\}$. Both programs are implemented in Matlab; the LP using linprog from Matlab's optimisation toolbox, and the ILP using milp downloaded from http://www.ie.ncsu.edu/kay/matlog, and written by Michael G. Kay. Our implementation of the SDP policy iteration algorithm is also run in Matlab.

2.3 Inflow probabilities

The relationship of runoff to rainfall is generally considered to be non-linear due to the influences of vegetation cover, infiltration rate, ground slope, rainfall intensity and soil moisture content. However, a linear relationship is a reasonable approximation for calculating runoff for a given catchment (FAO, 1991), particularly for an urban catchment with a high proportion of sealed surfaces. Figure 2 shows annual rainfall at a site (Parafield) in the catchment of the stormwater harvesting site. Initially we assume that annual rainfall less than 380 mm provides 0 units of inflow to the site, rainfall amounts between 380 and 580 mm provide 1 unit of inflow and rainfall amounts greater than 580 mm provide 2 units. Yearly inflow probabilities for stormwater (Table 3) are calculated from the histogram of annual rainfall.

Table 3. Inflow probabilities, (p_k), and amounts (in units of volume) for the 3 - state, annual time step system

rainfall (mm)	**inflow amount**	p_k
minimum - 380	0	0.2
380 - 580	1	0.5
580 - maximum	2	0.3

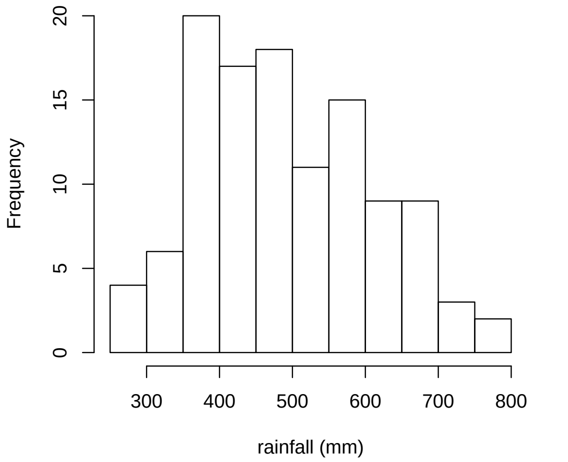

Figure 2. Annual rainfall at Parafield, Adelaide, 1885 - 1998

2.4 The stochastic dynamic program

The state space for the SDP algorithm is the content of the reservoir at the beginning of a time period $t = 1, \dots, T$. Decisions are defined as an intended release of $d \in \{0, 1, \dots, D\}$ units of water from the reservoir over the duration of the time step. If $d > k$ the ability to take $d - k$ units from the reservoir depends on the inflows to the reservoir.

The SDP algorithm uses transition matrices whose entries, $p_{kl}(d)$, which depend on a decision d, are the probability of moving between states of the state space, and a reward matrix whose entries, $r_{kl}(d)$, are the value obtained by making a particular transition under decision d. A transition matrix is specified for each decision and a corresponding reward matrix calculated for each decision. Let $k \in \{1, \dots, K\}$ represent the states of the system. For our initial model, $K = 3$. The policy iteration procedure is implemented in two parts - value determination and policy improvement (Howard, 1960). The policy iteration procedure assumes that T is sufficiently far in the future for a steady state to apply. For a given policy, total expected earnings over the remaining time steps at time t depends on the state, k, at time t, and is written as $v_t(k)$. For a given policy, total expected earnings is calculated recursively as

$$v_t(k) = \sum_l p_{kl}[r_{kl} + v_{t+1}(l)] \qquad \text{for } l = 1, \dots, L. \tag{2}$$

For large T, in the steady state,

$$v_t(k) = g + v_{t+1}(k) \tag{3}$$

where g is the expected return per period. The expected return is a constant, regardless of the current state, because the additional step is sufficiently far in the future to be independent of the current state (the ergodic property). Substituting (3) into (2) gives the set of equations making up the value determination step

$$g + v(k) = \sum_l p_{kl} r_{kl} + \sum_l p_{kl} v(l) \qquad k = 1, \dots, K \tag{4}$$

which are solved for g and $v(2)$ up to $v(K), v(1)$ being arbitrarily set at 0. The policy improvement step maximises over d for all states k

$$\sum_k p'_{kl}(d) r'_{kl}(d) + \sum_l p'_{kl}(d) v(l). \tag{5}$$

The algorithm starts with an arbitrary policy and continues until the policies produced on two successive iterations are identical.

We assume there are no losses of water from the reservoir other than intended withdrawals. Given the probability vector of inflows over a single time step $(p_0\, p_1\, p_2)$ for inflows of 0, 1 and 2 units of volume respectively, the transition matrix for decision 0 is

$$P(d_0) = \begin{bmatrix} p_0 & p_1 & p_2 \\ 0 & p_0 & p_1 + p_2 \\ 0 & 0 & 1 \end{bmatrix}$$

The transition matrix for decision 1 is

$$P(d_1) = \begin{bmatrix} p_0 + p_1 & p_2 & 0 \\ p_0 & p_1 & p_2 \\ 0 & p_0 & p_1 + p_2 \end{bmatrix}$$

The transition matrix for decision 2 is

$$P(d_2) = \begin{bmatrix} 1 & 0 & 0 \\ p_0 + p_1 & p_2 & 0 \\ p_0 & p_1 & p_2 \end{bmatrix}$$

The reward associated with a decision depends on the decision to take a number of units of stormwater from the reservoir, the probability of those units being available during the time period, and the allocations of water made from this and the other sources to meet demand. If a decision is taken, for example, to take 2 units of water while the current state of the system is 0, with probability p_0 the decision may result in 0 units of stormwater being taken from the reservoir, with probability p_1 1 unit is taken from the reservoir and with probability p_2 2 units are taken from the reservoir. The LP or the ILP, as appropriate, determine the optimal allocations for each state of the problem, given the supply, demand and salinity constraints. The profit, E_ν, from the optimum allocated given by the LP/ILP depends on the units of stormwater available (ν). The units of stormwater available depend on the state when a decision is made and the random input to the reservoir. These profits, weighted by the probability of taking the amount of stormwater from the reservoir under the decision, make up the reward matrix, $r(d)_{kl}$, for each decision. The reward matrix for decision 0 is

$$r(d_0) = \begin{bmatrix} E_0 & E_0 & E_0 \\ E_0 & E_0 & E_0 \\ E_0 & E_0 & E_0 \end{bmatrix}$$

The reward matrix for decision 1 is

$$r(d_1) = \begin{bmatrix} \frac{p_0}{(p_0+p_1)}E_0 + \frac{p_1}{(p_0+p_1)}E_1 & E_1 & E_1 \\ E_1 & E_1 & E_1 \\ E_1 & E_1 & E_1 \end{bmatrix}$$

The reward matrix for decision 2 is

$$r(d_2) = \begin{bmatrix} p_0E_0 + p_1E_1 + p_2E_2 & E_2 & E_2 \\ \frac{p_0}{(p_0+p_1)}E_1 + \frac{p_1}{(p_0+p_1)}E_2 & E_2 & E_2 \\ E_2 & E_2 & E_2 \end{bmatrix}$$

A policy is defined by specifying a decision for each state of the system. For example, one policy is to withdraw the current contents of the reservoir over the time step. The three decisions are: withdraw ν units of stormwater if the reservoir contains ν units ($\nu \in \{0, 1, 2\}$). The decisions that make up this policy can always be implemented. By using a vector to represent each state of the system in order and the elements of the vector to represent decisions, we can write the above policy as $[0, 1, 2]^T$. A contrasting policy would consist of decisions to withdraw 2 units of water from the reservoir over the time step, $[2, 2, 2]^T$. For this policy, decisions can only be implemented from states 0 and 1 if the inflow to the reservoir is 2 or 1 respectively.

3 POLICIES AND PROFITS

3.1 Initial model - 3-state, yearly time step

The optimal allocations found by the ILP for 0, 1 or 2 units of stormwater are shown in Table 4 below. As stormwater becomes available, use of the most expensive source of water - mains - is reduced. State 2 enables the blending of stormwater with recycled and mains water to meet the salinity conditions for urban use and thus the units of water supplied in that state increases to 14 units. By contrast, the allocations for the LP and 3-state system in Table 5 show a slight reduction in overall water supply as stormwater availability increases. This is a consequence of the cost of mains water exceeding the return for preferred demand.

The SDP finds an optimal policy for the ILP and EMV objective of $[0, 2, 2]^T$. This policy has transition matrix

$$P_{[0,2,2]} = \begin{bmatrix} 0.2 & 0.5 & 0.3 \\ 0.7 & 0.3 & 0 \\ 0.2 & 0.5 & 0.3 \end{bmatrix}$$

Table 4. Optimal allocations found by the ILP for each state for the 3-state, annual time step model

state of reservoir	source	allocations			total supply	profit
		wool	urban	council		
$\{0\}$	storm	0	0	0	12	18000
	recycled	0	1	3		
	mains	2	3	3		
$\{1\}$	storm	0	1	0	11	19300
	recycled	0	1	3		
	mains	2	1	3		
$\{2\}$	storm	0	2	0	14	22050
	recycled	0	2	3		
	mains	2	2	3		

Table 5. Optimal allocations found by the LP for each state for the 3-state, annual time step model

state of reservoir	source	allocations			total supply	profit
		wool	urban	council		
$\{0\}$	storm	0	0	0	14	19220
	recycled	0	1.6	3.2		
	mains	2	4.4	2.8		
$\{1\}$	storm	0.32	0.34	0.34	13.75	20812
	recycled	0.09	1.63	3.3		
	mains	1.59	3.78	2.37		
$\{2\}$	storm	0.51	0.8	0.69	12.75	22112
	recycled	0.14	1.48	3.38		
	mains	1 36	2.47	1.93		

which has equilibrium distribution $(0.4083, 0.4167, 0.1750)$ and long-term profit of \$20,167. Long-term profit is the product of the proportion of time the system spends in each state in the long-term - the equilibrium distribution - and a weighted sum of the rewards for the decision prescribed for each state. The weights are based on the probability distribution of inflows. For example, the long-term profit for the ILP, EMV objective, 3-state system given above is $0.4083 \times 18000 + 0.4167 \times (0.2 \times 19300 + 0.8 \times 22050) + 0.175 \times 22050 = 20167$.

The trend in supply for the LP is to provide the guaranteed amount of water to the wool processor, the maximum amount to the council and progressively reduce supply to the urban greywater network as more stormwater is available. The trend in use of sources is to make full use of any stormwater, use as much recycled water as possible - given the salinity restrictions, and reduce use of the most expensive mains water. The SDP finds an optimal policy for the LP of $[1, 2, 2]^T$. The transition matrix for policy $[1, 2, 2]^T$ is

$$P_{[1,2,2]} = \begin{bmatrix} 0.7 & 0.3 & 0 \\ 0.7 & 0.3 & 0 \\ 0.2 & 0.5 & 0.3 \end{bmatrix}$$

which has equilibrium distribution $(0.7, 0.3, 0)$ and long-term profit of \$20,901. With its freedom to blend fractional amounts of water and find an optimal solution over a continuum, the LP solution does produce a higher long-term profit as expected. However, the percentage increases in profit found in moving from the optimal ILP solution to the optimal LP solution are small, being 6.78, 7.83 and 0.28 % for each state respectively. The potentially better representation of the system and thus greater accuracy obtained by integer solutions may not be justified if they increase computational complexity, particularly if a higher dimensional representation of the system is used.

3.2 Policy change under CVaR objective

CVaR is calculated using the definition given in Section 1. Here, the worst case losses correspond to lowest profits and occur when inflow to the reservoir is 0. The probability of 0 inflow is 0.2 so we work with a $\text{CVaR}_{0.8}$ criterion, and a time horizon of one year.

The CVaR objective appears in the SDP algorithm in the value determination step as

$$h_t(i) = [(\text{CVaR}_\alpha)_i + \sum_j p_{ij} h_{t+1}(j)] \qquad (6)$$

The optimal policy with $\text{CVaR}_{0.8}$ for both the linear and integer programs, is $[0, 0, 2]^T$. The policy shows the potential value of holding water in storage to reduce the likelihood of having to obtain water from a more expensive source in a later period. This policy has an equilibrium distribution of (0.1231, 0.3846, 0.4923), long-term profit of $20,644 for the LP and $19,994 for the ILP. This is the expected approach from using a CVaR objective as the risk measure is conservative - it favours avoiding the conditions which could generate the worst loss.

3.3 Increasing resolution - the 5-state model

We increase the state space of the reservoir and find a vector of inflow probabilities for amounts of 0, 1, 2, 3 or 4 units of stormwater of [0.0915, 0.3384, 0.2652, 0.2195, 0.0854] respectively. Salinity constraints are as before. Cost and return per unit (2) are halved to enable direct comparison with the 3-state model.

Optimal water allocations for each state, total supply and profit for the LP are shown in Table 7, as are equivalent results for the ILP in Table 8. The optimal policy for the LP with EMV objective is $[2,2,2,3,4]^T$. The equilibrium distribution for this policy is (0.4568, 0.2647, 0.1710, 0.0837, 0.0238) and long-term profit is $20,698. The optimal policy for the ILP with EMV objective is $[0,0,3,3,3]^T$. This policy has equilibrium distribution (0.1523, 0.2267, 0.2795, 0.1989, 0.1426) and long-term profit $20,366.

The probability of the worst outcome - having 0 units of stormwater available - is 0.09 for the 5-state model, thus α for CVaR is now 0.91. The $\text{CVaR}_{0.91}$ objective gives an optimal policy of $[0,0,0,3,3]^T$ for the ILP. The equilibrium distribution for this policy is (0.0272, 0.1417, 0.2539, 0.2702, 0.3070) and long-term profit $20,284. Applying the $\text{CVaR}_{0.91}$ objective to the LP gives an optimal policy of $[0,1,2,3,4]^T$. The equilibrium distribution for this policy reproduces the inflow probabilities and long-term profit is $20,678. Under policy $[0,1,2,3,4]^T$ managers take only the amount of stormwater currently held in the reservoir.

Table 6. Availability and demand for sources and sinks for the 5-state problem. Availability, firm and preferred demand are in units of volume. Stormwater availability is a random value $\in \{0, 1, 2, 3, 4\}$.

source:	**mains**	**recycled**	**storm**
availability	up to 40	up to 10	$\{0, 1, 2, 3, 4\}$
sink:	**wool**	**urban**	**council**
firm demand	4	6	6
pref. demand	6	12	12

Table 7. Optimal allocations found by the LP for each state for the 5-state, annual time step model

state	source	allocations wool	urban	council	total supply	profit
{0}	storm	0	0	0	28	19220
	recycled	0	3.20	6.40		
	mains	4	8.80	5.60		
{1}	storm	0.28	0.34	0.39	28	20065
	recycled	0.08	3.29	6.50		
	mains	3.65	8.37	5.11		
{2}	storm	0.59	0.71	0.70	27.5	20812
	recycled	0.16	3.26	6.59		
	mains	3.25	7.54	4.71		
{3}	storm	0.72	1.22	1.06	26.5	21462
	recycled	0.19	3.13	6.68		
	mains	3.09	6.15	4.26		
{4}	storm	0.77	1.88	1.35	25.5	22112
	recycled	0.21	3.03	6.76		
	mains	3.02	4.59	3.89		

Table 8. Optimal allocations found by the integer program for each state for the 5-state, annual time step model

state	source	allocations wool	urban	council	total supply	profit
{0}	storm	0	0	0	28	18725
	recycled	0	3	6		
	mains	4	9	6		
{1}	storm	0	0	1	27	19375
	recycled	0	3	6		
	mains	4	9	4		
{2}	storm	0	1	1	26	20025
	recycled	0	3	6		
	mains	4	7	4		
{3}	storm	0	0	3	28	21425
	recycled	0	3	7		
	mains	4	9	2		
{4}	storm	0	3	1	27	22075
	recycled	0	4	6		
	mains	4	5	4		

3.4 Increasing resolution - 3-state 2-season model

The one year decision period of the initial model may be insufficiently detailed for practical management. Adelaide has a typical Mediterranean rainfall pattern and so it seems natural to divide the year into 2 seasons of drier and wetter periods. Doing so against average monthly rainfall over the 107 year record produces 2 seasons of uneven length, the mostly dry period (summer) of 5 calendar months from November to March and the wetter period (winter) of 7 months from April to October (Figure 3).

Representing stormwater storage with 3 states and specifying the amounts of stormwater availability as 0 for rainfall less than 140 mm for the season, 1 for rainfall between 140 and 240 mm, and 2 for rainfall above 240 mm, we find transition probabilities from the histograms of Figure 3. A probability vector for summer inflows of 0, 1 or 2 units of stormwater is $(0.74, 0.23, 0.03)^T$. Formally the SDP is now set up with 6 states $s0, s1, s2, w0, w1$ and $w2$ representing the seasons and states of the reservoir at the beginning of the time step.

The optimal policy for the summer months found by the LP and EMV is $[1, 2, 1]^T$, while that for the ILP and EMV is $[0, 0, 2]^T$. The equilibrium distributions for these optimal policies are (0.97, 0.02, 0.01) for policy $[1, 2, 1]^T$ and (0.39, 0.48, 0.13) for $[0, 0, 2]^T$. Note that the probability of seeing 2 units of inflow of stormwater in a single time step is only 0.02, thus the optimal ILP policy has essentially a strategy of conserving a single unit of water when it is available in one summer, and carrying it over to the next summer when a second unit may inflow to the reservoir. This is not a practical management policy. Long-term profit for policy $[1, 2, 1]^T$ is \$19,362 and \$18,526 for $[0, 0, 2]^T$.

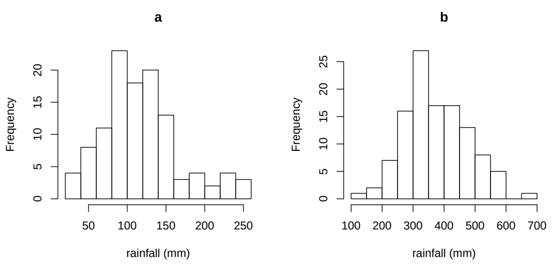

Figure 3. Parafield rainfall for a. summer, defined as the 5 months November to March and b. winter, defined as the 7 months April to October.

A probability vector for winter inflows with stormwater availability as specified for summer is $(0.01, 0.06, 0.93)^T$. The optimal policy for the winter period found by the LP and EMV is $[1, 2, 2]^T$, the same as found for the entire year, while that for the ILP and EMV is $[0, 2, 2]^T$. The equilibrium distributions for these optimal policies are (0.07, 0.93, 0) for policy $[1, 2, 2]^T$ and (0.04, 0.46, 0.50) for $[0, 2, 2]^T$. As for a whole year decision period, the optimal policy for use of stormwater in winter has the effect of managing the reservoir so that, in the long-term, it is never full. Here, this policy has an economic imperative but the policy is also environmentally friendly, providing capacity to intercept stormwater which would bypass a full reservoir and be released to nearby estuaries with deleterious effects on near-shore habitat. Long-term profit for policy $[1, 2, 2]^T$ is \$21,947 and \$21,812 for $[0, 2, 2]^T$.

4 CONCLUSION

Our integer/linear stochastic dynamic programming model generated policies for the optimal management of stormwater in a blending problem. Policies found by EMV criterion were usually found to take stormwater in amounts above those held in storage at the beginning of the time step, while the CVaR criterion produced policies that conserved water currently held in storage for some states. Policies found by LP and by ILP under the EMV criterion were different and long-term profits differed slightly - the difference decreased when resolution of the state space was increased. Increasing the resolution of the time variable demonstrated that seasonal differences lead to differing optimal policies. To make a more realistic allowance for the stochastic variation in stormwater throughout the year we will need to discretise time at a finer scale and hence the volume discretisation must increase.

5 REFERENCES

www.fao.org/docrep/U3160E/u3160e05, accessed 30/07/07.

T.W. Archibald, K.I.M. McKinnon, and L.C. Thomas. Modeling the operation of multireservoir systems using decomposition and stochastic dynamic programming. *Naval Research Logistics*, 53:217–225, 2006.

J. Cabero, Á. Baíllo, S. Cerisola, M. Ventosa, García-Alcalde, F. Perán, and G. Relaño. A medium-term integrated risk management model for a hydrothermal generation company. *IEEE Transactions on Power Systems*, 20:1379–1388, 2005.

R.A. Howard. *Dynamic programming and Markov processes*. The MIT Press, Cambridge, Mass., 1960.

R.T. Rockafellar and S. Uryasev. Conditional value-at-risk for general loss distributions. *Journal of Banking and Finance*, 26:1443–1471, 2002.