Water property rights in rivers with large environmental water holders

Neal Hughes

Australian National University

May 4, 2015

Abstract

This paper considers the design of water property rights in river systems — such as the Australian Murray-Darling Basin (MDB) — where large Environmental Water Holders (EWHs) are active. In particular, it considers the definition of rights to dam storage capacity and to physical river flows. We present a decentralised model of a regulated river system involving a large number of consumptive water users (i.e., farmers) and a single large EWH. These entities each make private water storage and trade decisions, subject to the prevailing property right rules, so as to maximize their own objectives (i.e., farm profits or environmental benefits). We specify broad parameter ranges — reflective of rivers in the MDB — and present the results of a large number of model runs. We find the ideal approach to storage rights is the 'capacity sharing' model advocated by Dudley and Musgrave (1988). In contrast, poorly specified storage rights can lead to large external effects on consumptive users. Further, we find that priority flow rights outperform simple proportional flow rights. Low priority water rights are found to be a good match for the demands of EWHs, helping to minimize their exposure to market transaction costs.

1 Introduction

Efforts to secure environmental flows are occurring in many of the worlds heavily regulated rivers. Frequently, this requires governments to become participants in water property right systems. That is, governments acquire water rights from consumptive users (through 'buyback' or otherwise) then 'use' the water allocations they receive to achieve environmental objectives

In Australia, the government has committed over $13 billion to acquiring water rights in the MDB. These rights are held by the Commonwealth Environmental Water Holder (CEWH). The CEWH joins many other smaller Environmental Water Holders (EWH) already active in the basin[1].

The participation of EWH's in water markets raises a number of policy questions. Water property rights have evolved over a long period of time, to satisfy the requirements of consumptive users (i.e., irrigation farmers). EWH's have very different patterns of water demand and face very different incentives to existing users.

The introduction of an EWH — with payoffs defined over in-stream flows as opposed to extractive use — raises obvious externality problems. Consumptive users affect in-stream flows, through their water deliveries, return flows, and storage reserves (via spills). Similarly, EWH decisions may affect the availability of water for consumptive users. As Brennan (2011) explains, these effects may depend on the nature of storage rights:

> An increase in [environmental flows] will have implications for irrigation reliability. The extent to which irrigators can mitigate these reliability impacts depends on the rules governing storage and how capacity is divided between the irrigation industry and the environmental water-holder. Different rules of storage management might result in different outcomes for the environment as well as for irrigation reliability (Brennan 2011; pp. 309)

An additional concern is market power since some EWHs — such as the CEWH — can be much larger than irrigators. Large water holders are more likely to take into account their effect on aggregate variables such as prices, storage volumes and spills, as Brennan (2011) argues:

> The Water Act permits the ...[CEWH]... to trade on the seasonal market, if it is beneficial for achieving environmental outcomes. But the presence of a large player in the market might have a significant impact on

[1]Including the Victorian Environmental Water Holder and NSW RiverBank

> prices. Rules governing the participation of the environmental water-holder in the seasonal market require serious consideration (Brennan 2011; pp. 311)

This paper presents a computational model of a regulated river system involving a large number of consumptive water users (i.e., farmers) and a single large EWH. This paper adds to the literature which models the behaviour of EWHs in water markets, in particular their water trade and water storage decisions (for example Grafton et al. 2011, Heaney et al. 2011, Beare et al. 2006, Kirby et al. 2006).

However, the main contribution is a comparison of water property right systems. To achieve this we develop a decentralised version of our model in which each agent (including all the farmers and the EWH) make private water storage and trade decisions subject to prevailing property right rules, so as to maximize their own objectives (i.e., farm profits or environmental benefits).

Given the presence of externalities the problem is a stochastic game: each agent faces a Markov decision problem where the payoffs and state transitions depend on the other agents. We solve this model numerically via a novel 'multi-agent learning' algorithm. Effectively this allows us to populate the model with 'intelligent' (i.e., near optimal) selfish agents.

Armed with this model we ask the question: which form of water property rights is ideal in the presence of a large EWH. In particular, we compare various approaches to dam storage rights (including the capacity sharing model of Dudley and Musgrave 1988) and contrast priority and proportional type rights to physical river flows.

This paper proceeds as follows. First, we review of the economic literature attempting to model EWHs. Next we define the planner's version of the model and present some illustrative results. We then define the decentralised version of the model and present the results of a series of policy scenarios.

2 Literature

There are small number of economic studies modeling environmental flows in regulated rivers (Dudley et al. 1998, Beare et al. 2006, Kirby et al. 2006, Heaney et al. 2011, Grafton et al. 2011). Existing studies focus more on the role of the government in acquiring, using and trading water for environmental purposes, than on externalities. Typically, researchers have used single agent models or employed simplifying assumptions: either holding the behaviour of some of the agents fixed or limiting the feedback between the agents.

Dudley et al. (1998) present a complex model of the Barker-Barambah catchment in southern Queensland. This model builds on Dudley (1988)'s earlier work, combining farm level models with a river flow model and a set of environmental objectives[2]. A major focus is estimating trade-off curves comparing environmental and irrigation benefits for varying allocations of water rights to the environment.

Beare et al. (2006) consider the release decisions of an EWH concerned with over-bank flow events, using a model of the Murrumbidgee river. The model assumes a high flow event with a target size, timing, duration and frequency. The EWHs objective is to minimise penalties for failing to meet these targets and water resource costs. Beare et al. (2006) suggest EWHs should hold large volumes of low priority water rights and then sell unneeded allocations back to irrigators.

Heaney et al. (2011) present a model of the Goulburn region and again focus on flow inter-arrival times. In particular, Heaney et al. (2011) show that economic costs are sensitive to small increases in the 'reliability' of (the probability of achieving) target inter-arrival times. They go on to demonstrate the welfare gains from short term EWH trading (i.e., the spot market) and carryover (i.e. inter-year storage reserves).

Grafton et al. (2011) present a social planners SDP model of the Murray river. Grafton et al. (2011) adopt an environmental objective similar to that of Heaney et al. (2011) with an increasing penalty for delaying high flow events beyond a target inter-arrival time. Grafton et al. (2011) present stylised optimal EWH release rules and estimate the welfare gains from optimal versus observed historical environmental flows for the Murray.

There are few analytical studies on EWHs in regulated rivers. Truong (2011) considers the effects on consumptive users of an exogenous reduction in the storage capacity within a theoretical model. While the paper is concerned with the external effects of an EWH the model does not include any representation of environmental objectives.

[2]Similar to Dudley (1988) the model has nested structure - with distinct environmental and irrigation problems, each over different time steps.

3 The planner's model

This paper is concerned with the allocation of water within an abstract regulated river system (figure 1).

Figure 1: An abstract regulated river system

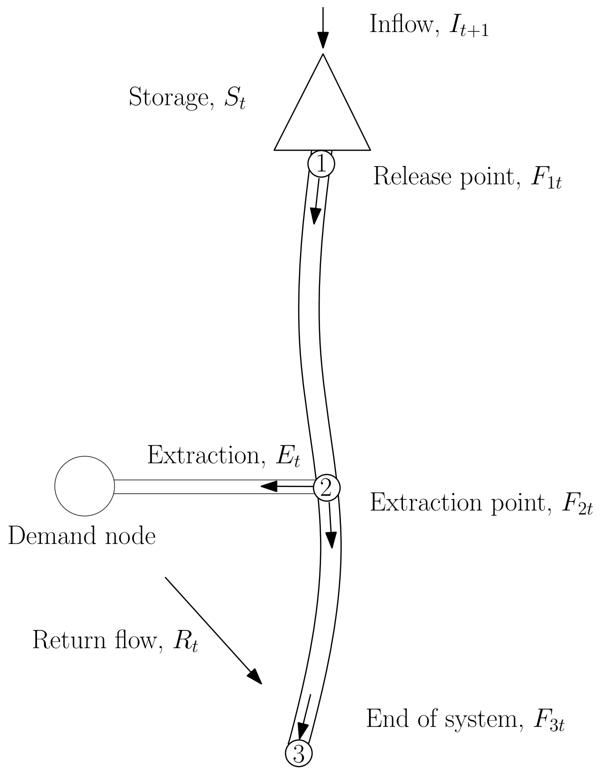

The model involves a single reservoir with fixed capacity receiving stochastic inflow. Water released from storage can be extracted and supplied to a single demand node (i.e. irrigation area), populated with a large number of heterogeneous water users (i.e. irrigation farmers) . In addition, we can have entities that value in-stream flows, such as environmental water holders.

3.1 Inflow

The model adopts a bi-annual time scale, dividing the year into a 'winter' (April-September) and 'summer' (October-March) season. M_t indicates the season: 0 is summer and 1 winter.

C_t is an unobservable climate state variable (i.e., annual inflow). C_t follow an *annual* AR-1 process with gamma shocks

$$C_{t+1} = \begin{cases} \rho_C C_t + \epsilon_{t+1} & \text{if } M_t = 1 \\ C_t & \text{if } M_t = 0 \end{cases}$$

$$0 < \rho_C < 1$$

$$\epsilon_{t+1} \sim \Gamma(k_C, \theta_C)$$

Seasonal inflows I_{t+1} are then defined

$$I_{t+1} = \begin{cases} \omega_t C_t & \text{if } M_t = 0 \\ (1 - \omega_t) C_t & \text{if } M_t = 1 \end{cases}$$

$$\omega_t \sim N_{\omega_a}^{\omega_b}(\mu_\omega, \sigma_\omega)$$

Where N_a^b denotes the truncated normal distribution with support $[a, b]$.

3.2 Storage

The storage transition rule is

$$S_{t+1} = \min\{\max\{S_t - W_t - L_t + I_{t+1},\, 0\},\, K\}$$

$$0 \leq W_t \leq S_t$$

Where S_t is the storage volume, W_t the storage withdrawal (release), K the fixed storage capacity and L_t is evaporation loss. L_t is a standard concave function of storage volume (following Lund 2006):

$$L_t = \delta_{0t}.\alpha(S_t)^{2/3}$$

Storage spills Z_{t+1} are then defined as

$$Z_{t+1} = \max\{0,\, I_{t+1} - (K - S_t + W_t + L_t)\}$$

3.3 River flow

River flow volumes F_{jt} during period t are defined for each node j (as indicated in figure 1):

$$F_{1t} = W_t + Z_t$$

$$F_{2t} = F_{1t} - \mathcal{L}_1(F_{1t}) - E_t$$

$$F_{3t} = F_{2t} - \mathcal{L}_2(F_{2t}) + R(E_t)$$

where E_t is river extraction, R is a return flow function and $\mathcal{L}_1, \mathcal{L}_2$ are delivery loss functions:

$$\mathcal{L}_1(F_{1t}) = \min\{F_{1t}, \delta_{at}\}$$

$$\mathcal{L}_2(F_{2t}) = \min\{F_{2t}, \delta_{at}\}$$

Return flows are assumed to be a fixed proportion of extraction

$$R(E_t) = \delta_R E_t$$

3.4 Extraction

We assume spills are unavailable for consumptive use, so that extraction is constrained by releases less losses

$$E_t \leq W_t - \mathcal{L}_1(W_t)$$

Total water use $Q_t = \sum_{i=1}^{n} q_{it}$ is constrained by extraction less delivery losses

$$Q_t \leq \max\{(1 - \delta_{Eb})E_t - \delta_{Ea}, 0\}$$

3.5 Consumptive demand

There n consumptive users $\mathcal{U} = \{i \in 1, ..., n\}$. The set of users $\mathcal{U}$ is partitioned into a low reliability group (e.g., broadacre) $\mathcal{U}_{low} = \{i \in 1, 2, ..., n_{low}\}$ and a high-reliability group (e.g., horticulture) $\mathcal{U}_{high} = \mathcal{U}_{low}^{\complement}$. $h \in (high, low)$ indicates membership to these sets.

All consumptive use occurs in the summer period (the irrigation season). Each group has a quadratic relationship between between water use q_{it} and profit:

$$\pi_{ht}(q_{it}, I_{it}, e_{it}) = \mathscr{A}_h.e_{it}(\theta_{h0} + \theta_{h1}\tilde{q}_{it} + \theta_{h2}\tilde{q}_{it}^2 + \theta_{h3}\tilde{I}_t + \theta_{h4}\tilde{I}_t^2 + \theta_{h5}\tilde{I}_t.\tilde{q}_{it})$$

Here π_h is the profit function of users in class h, $\mathscr{A}_h$ is the fixed land area for each user in class h, $\tilde{q}_{it} = q_{it}/\mathscr{A}_h$ is water use per unit of land and $\tilde{I}_t = I_t/E[I_t]$ is a proxy for local climate conditions. Finally, e_{it} is a user specific productivity shock following an (annual) AR(1) process.

3.6 Environmental demand

The EWHs objective is broadly to minimise the deviation between 'natural' and actual river flows, similar to the approach of Dudley et al. (1998).

Here, natural river flows $\tilde{F}_{jt}$ are those that would have prevailed in the absence of regulation: where $F_{1t} = I_t$. We define ΔF_{jt} as a measure of the deviation between $\tilde{F}_{jt}$ and F_{jt}[3]

$$\Delta F_{jt} = \begin{cases} \min\left\{\left(\frac{\tilde{F}_{jt} - F_{jt}}{\tilde{F}_{jt}}\right)^2, 1\right\} & \text{if } \tilde{F}_{jt} > 0 \\ \mathbb{1}_{F_{jt} > 0} & \text{if } \tilde{F}_{jt} = 0 \end{cases}$$

The environmental benefits (in dollars) are then

$$B(.) = b_\$ e_{0t} G_I(I_t)(1 - \sum_{j=1}^{3} b_j \Delta F_{jt})$$

$$\sum_{j=1}^{3} b_j = 0, 0 < b_j < 1, b_\$ > 0$$

$$e_{0t} \sim N_0^2[1, \sigma_{e0}]$$

$$G_I(I_t) = \Pr(I \leq I_t)$$

Here the b_j parameters determine the relative importance of each flow node and $b_\$$ determines the importance of the environment relative to profits. e_{0t} reflects exogenous variation in demand for environmental flows (all factors other than river flows which influence ecological condition).

Following Dudley et al. (1998) environmental payoffs are weighted by the Cumulative Distribution Function (CDF) for inflows G_I, which increases the incentive to release water in high flow years[4].

3.7 The planner's problem

The planner's problem is to determine storage releases W_t, extraction E_t and water use q_{it} each period — conditional on state variables S_t, I_t, e_{it}, M_t — to maximise social welfare

[3] Here ΔF_{jt} is just the squared percentage deviation, adjusted for two special cases. One, where actual river flows are more than double natural flows. Two where natural flows are zero.

[4] In the absence of these weights, the model tends to focus too much on low flow years, given the lower opportunity costs (i.e. release volumes).

$$\max_{\{q_{it},W_t,E_t\}_{t=0}^{\infty}} E\left\{\sum_{t=0}^{\infty}\beta^t\left(\sum_{i=1}^{n}\pi_{ht}(q_{it},\tilde{I}_{it},e_{it})+B(.)\right)\right\}$$

subject to the constraints detailed above.

4 Parameterisation

To maintain generality, parameter ranges are specified rather than point estimates. Complete detail on the parameterisation is provided in Hughes (2015).

Supply side parameters are based on statistics for 22 storages in the MDB (all those greater than 50 GL in capacity). A data set on these storages was compiled from various sources including NWC (2011), ANCOLD (2013) and BOM (2013). Parameter distributions are assumed uniform over the 15th to the 85th percentiles of our data set.

Demand side parameters are based primarily on an econometric analysis of irrigation farms in the southern MDB, drawing on ABARES survey of irrigation farms (Ashton and Oliver 2012). High reliability water demand is estimated using a sample of southern MDB grape farms and low reliability a sample of southern MDB broadacre farms.

For the environmental demands we assume $b_2 = 0$ and

$$b_1 \sim U[0, 0.6]$$
$$\frac{b_\$}{\bar{C}} \sim U[30, 130]$$
$$\sigma_{e0} \sim U[0, 1]$$

$b_\$$ is set so the reduction in extraction under the central case parameters is comparable with that proposed under the basin plan (25.6 per cent under the planner's solution)

Storage capacity K is the numéraire in parameterisation and is fixed at 1000 GL. Key parameter ranges are shown in table 1.

Table 1: Selected parameter ranges

	Min	Central case	Max
$E[C_t]/K$	0.23	0.71	1.18
$\text{Var}[C_t]^{0.5}/E[C_t]$	0.40	0.70	1.00
ρ_C	0.20	0.25	0.30
μ_ω	0.55	0.65	0.75
σ_ω	0.09	0.105	0.12
$\alpha K^{2/3}/K$	0.03	0.09	0.15
δ_R	0.0	0.10	0.2
$\delta_{Ea}/E[C_t]$	0.0	0.05	0.1
δ_{Eb}	0.1	0.2	0.3
n	100	100	100
n_{low}	30	50	70
n_{high}	30	50	70
ρ_e	0.30	0.40	0.50
σ_η	0.10	0.15	0.20
b_1	0.0	0.3	0.6
$b_\$/E[C_t]$	30	80	130
σ_{e0}	0	0.5	1

5 The planner's solution

We solve the planner's problem using a reinforcement learning (aka approximate dynamic programming) algorithm: fitted Q-V iteration (see Hughes 2015). We use the spot market equilibrium conditions to solve the use allocation problem. The planner then has a MDP with one policy variable W_t and four state variables: S_t, I_t, e_{0t}, M_t[5].

We consider two solutions to the planners problem: a 'consumptive' case where environmental benefits are ignored (where $b_\$$ and an 'optimal' case where both are considered. Annual results for the central case parameters are summarised in tables 3 to 8.

As would be expected, the optimal scenario achieves higher social welfare, lower profits and greater environmental benefits relative to the consumptive case. The optimal scenario leads to a \$10.7m loss of mean profit per year and a \$17.8m gain in mean environmental benefits (for a net gain of \$7.1m).

The optimal scenario leads to a mean reduction in extraction of 145 GL or around 26 per cent. Mean storage and withdrawals show little change, but the variance of withdrawals does increase significantly.

Table 2: Social welfare $\sum_{i=0}^{n} u_{it}$ (\$M)

	Mean	SD	2.5th	25th	75th	97.5th
Consumptive	209.92	29.94	143.90	196.41	225.54	265.45
Optimal	217.00	39.88	146.91	193.05	242.59	296.20

Table 3: Consumptive user profits $\sum_{i=1}^{n} u_{it}$ (\$M)

	Mean	SD	2.5th	25th	75th	97.5th
Consumptive	193.93	21.55	140.30	186.96	207.57	218.59
Optimal	183.19	20.18	142.50	173.76	196.06	211.64

Table 4: Environmental benefits u_{0t} (\$M)

	Mean	SD	2.5th	25th	75th	97.5th
Consumptive	15.99	17.31	0.32	5.52	19.17	68.29
Optimal	33.81	30.72	1.38	7.76	52.54	108.81

[5]The planner ignores the user productivity shocks since with 100 users they have close to no aggregate effect. Here we also assume the planner conditions only on the latest inflow I_t, ignoring I_{t-1} which in this model has some relevance for forecasting I_{t+1}.

Table 5: Storage S_t (GL)

	Mean	SD	2.5th	25th	75th	97.5th
Consumptive	595.45	266.80	131.61	364.40	823.06	1,000.00
Optimal	588.18	269.92	133.92	352.77	827.04	1,000.00

Table 6: Withdrawal W_t (GL)

	Mean	SD	2.5th	25th	75th	97.5th
Consumptive	571.79	191.81	158.96	417.07	735.54	818.50
Optimal	592.93	235.63	175.53	421.09	776.20	1,030.16

Table 7: Extraction E_t (GL)

	Mean	SD	2.5th	25th	75th	97.5th
Consumptive	563.57	191.81	150.74	408.85	727.31	810.28
Optimal	418.88	149.79	155.17	301.43	534.96	688.52

Table 8: Shadow price P_t (GL)

	Mean	SD	2.5th	25th	75th	97.5th
Consumptive	125.08	187.51	0.00	52.40	142.34	633.50
Optimal	160.73	188.97	30.85	98.78	163.98	568.96

5.1 River flows

Figure 2 compares the natural river flow distribution with those of the optimal and the consumptive scenarios — in the form of *duration curves* [6]. Here we see how environmental flows offset changes to natural flow regimes caused by consumptive use. In particular, environmental flows result in: an increase in mean river flow (particularly downstream), an increase in the volatility of river flow and an increase in winter flow (in turn for a decrease in summer flow).

Figure 2 shows a significant increase in the frequency of small and medium flow events, but little to no change in the frequency of large flood events (e.g., greater than 1000 GL). This emphasis on low and medium flows is broadly consistent with environmental flow planning in the MDB. In addition to the high opportunity costs, the creation of large flow events has the potential to cause flood damage in practice.

[6] The duration curve shows the probability of *exceeding* a flow of a given magnitude: $1 - G_{F_{jt}}(.)$ where $G_F(.)$ is the CDF of F

Figure 2: River flow duration curves, summer

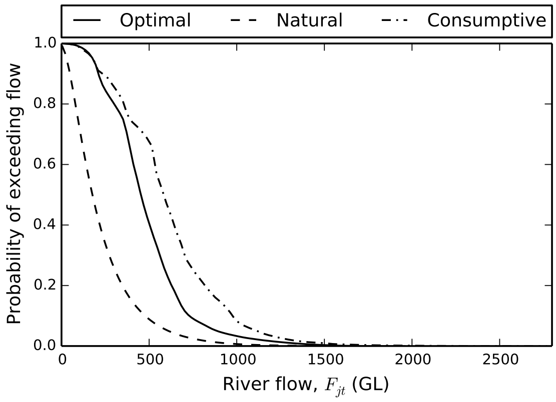

(a) Summer upstream, F_{1t}

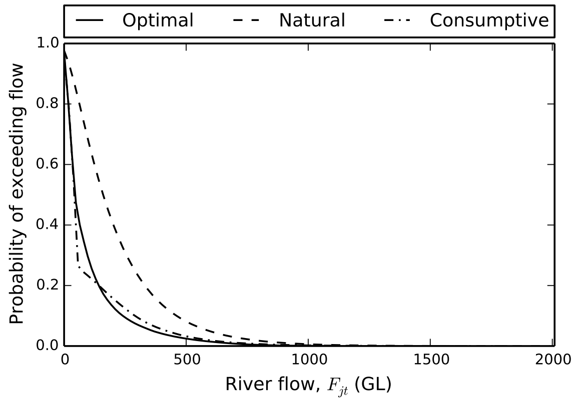

(b) Summer downstream, F_{3t}

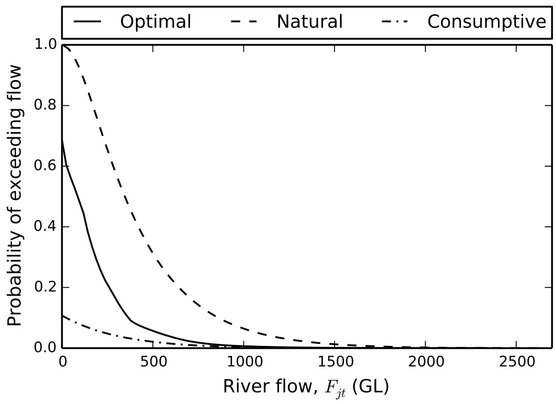

(c) Winter upstream, F_{1t}

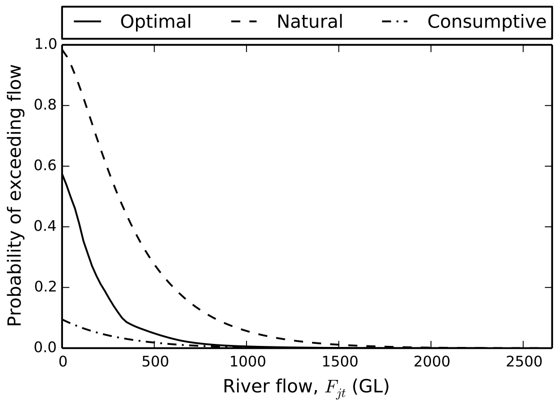

(d) Winter downstream, F_{3t}

6 The decentralised model

In the decentralised model, we have $n+1$ water right holders: n consumptive users $i = 1$ to n and one EWH $i = 0$. From a water accounting perspective the EWH is treated identically to other users.

6.1 Water property rights

Under the general water rights framework each user controls their own 'water account'. Users account balances s_{it} are credited with a share of inflow and debited for user withdrawals w_{it}, such that

$$\sum_{i=0}^{n} s_{it} = S_t$$
$$w_{it} \leq s_{it}$$

Further, storage releases W_t are a function of total user withdrawals $\sum_{i=0}^{N} w_{it}$ (see Hughes 2015).

To illustrate consider a capacity sharing (CS) scenario. Here user accounts are updated according to

$$s_{it+1} = \min\{\max\{s_{it} - w_{it} - l_{it} + \lambda_i I_{t+1} + x_{it+1}, 0\}, \lambda_i K\}$$
$$l_{it} = \left(\frac{s_{it}}{S_t}\right) L_t$$
$$\sum_{i=0}^{N} \lambda_i = 1$$

where λ_i is user i's share of inflow and capacity and x_{it+1} are storage 'externalities'.

Intuitively x_{it} are account reconciliations, which ensure the total account balance $\sum_{i=1}^{n} s_{it}$ matches the physical storage volume S_t. Under CS, x_{it} include 'internal spills': where a user's account reaches its limit and excess inflow is forfeited to other users. As such, x_{it} is a complex function of the account balances and withdrawals of all other users. Interaction between user accounts also occurs through the loss deductions l_{it}.

The CS scenario assumes inflows are shared proportionally, in practice we often have priority classes. In this paper we consider a two priority class system: high reliability users have priority over low reliability users. Formally, we replace $\lambda_i I_{t+1}$ above with

$$\min\{\lambda_i I_{t+1}, \lambda_i K\} \quad \text{if } i \in \mathcal{U}_{high}$$

$$\lambda_i(\max\{I_{t+1} - \sum_{\mathcal{U}_{low}} \lambda_i K, 0\}) \quad \text{if } i \in \mathcal{U}_{low}$$

We denote capacity sharing with priority inflow rights as scenario CS-HL.

6.2 Water share endowments

55 assume fixed water share endowments λ_i. The EWH's share λ_0 is based on the percentage change in extraction between the planner's optimal and consumptive scenarios $\hat{\lambda}_0$ (i.e., 25.6 per cent in the central case). The parameter distribution for λ_0 is then

$$\lambda_0 = N_0^1(\hat{\lambda}_0, 0.05)$$

User inflow shares are governed by a single parameter Λ_{high}: the total water share of high reliability users. The parameter distribution for Λ_{high} is

$$\Lambda_{high} = N_0^1(\hat{\Lambda}_{high}, 0.05)$$

where Λ_{high} is an estimate of the optimal mix of high/low priority rights conditional on n_{high} (for more detail see Hughes 2015).

Finally, we assume the EWH is a low priority user.

6.3 The spot market

All right holders receive water 'allocations' a_{it} in 'demand node' units [7]

$$a_{it} = w_{it}(1 - \delta_{Eb})$$

All users (including the EWH) can then participate in the spot market. In summer, the EWH can choose to sell all or a portion of their allocations to consumptive users or to purchase user allocations — effectively reducing extraction and increasing river flows at nodes 2 and 3. In winter, the EWH can purchase allocations from consumptive users.

[7] Fixed losses are 'socialised' (shared in proportion to inflow shares λ_i via the account reconciliation process).

We apply a transaction cost of $\tau/2$ to both sellers and buyers[8]. The users then have payoffs u_{it}

$$u_{it} = \begin{cases} \Pi_i + (P_t - \tau/2)(a_{it} - q_{it}) & \text{if } a_{it} - q_{it} \geq 0 \\ \Pi_i + (P_t + \tau/2)(a_{it} - q_{it}) & \text{if } a_{it} - q_{it} < 0 \end{cases}$$

where Π_i is the relevant benefit function: $\pi_{ht}(.)$ for farmers and $B(.)$ for the EWH and P_t is the market price for water.

In this context q_{0t} is 'environmental water consumption': storage releases that are not extracted, such that in summer

$$E_t = \delta_{Ea} + (\sum_{i=1}^{n} q_{it} - q_{0t})/(1 - \delta_{Eb})$$

In winter consumptive use $\sum_{i=1}^{n} q_{it}$ and extraction E_t are both zero, but users can still sell allocations to the EWH.

In Hughes (2015) we derive the spot market equilibrium conditions (the user and EWH water demand functions). Given quadratic benefit functions the demand functions are linear in the parameters. The spot market can then be solved independent of the storage (i.e., withdrawal) problems.

6.4 Users' problems

The users' problem is to maximise private benefits u_{it} by choosing w_{it}, q_{it}

$$\max_{\{q_{it}, w_{it}\}_{t=0}^{\infty}} E\left\{\sum_{t=0}^{\infty} \beta^t u_{it}\right\}$$

With q_{it} determined by spot market equilibrium, the users' problems have one policy variable w_{it} and five state variables $s_{it}, S_t, \tilde{I}_t, e_{it}, M_t$. Here we assume users condition only on the aggregate storage volume and are 'oblivious' to other users' account balances and productivity levels.

6.5 Environmental manager's problem

Similarly, the EWH's problem is to maximise u_{0t} by choosing w_{0t}, q_{0t}

[8] We assume a uniform distribution for τ over the range \$10 to \$100 per ML.

$$\max_{\{q_{0t},w_{0t}\}_{t=0}^{\infty}} E\left\{\sum_{t=0}^{\infty}\beta^t u_{0t}\right\}$$

The EWH's objective includes both environmental benefits and net trade proceeds / costs. As is typical in the literature we also apply a budget balancing constraint

$$\sum_{t=0}^{\infty} P_t\left(a_{0t} - q_{0t}\right) = 0$$

which prevents the EWH from accumulating a cash surplus (or deficit) in the long run — all trade proceeds must eventually be committed to environmental flows. We apply this constraint indirectly by varying the effective price faced by the EWH in the spot market (see Hughes 2015)[9].

With q_{0t} determined in the spot market, the EWHs problem, has one policy variable w_{0t} and five state variables $s_{it}, S_t, \tilde{I}_t, e_{0t}, M_t$.

6.6 Policy scenarios

Below we define our policy scenarios. Hughes (2015) provides more detail on each of these options and how they relate to property right systems in the Australian MDB and the western US.

6.6.1 Capacity sharing - CS

Scenarios CS, CS-HL are as defined above. CS reflects the capacity sharing approach to water property rights advocated by (Dudley and Musgrave 1988). Hughes (2015) shows that with no in-stream demands CS results generally outperform alternatives (SWA, NS, OA) and achieves aggregate welfare and storage volumes similar to a planner's outcome.

6.6.2 Spill forfeit rules ('spillable water accounts') - SWA

Spill forfeit rules are a common alternative to storage capacity rights. Here there are no limits on storage account volumes. However in the event of a physical storage spill, users are subject to deductions in proportion to their account volumes

[9]We adopt this approach mostly for computational reasons. Explicitly including the budget constraint would add both a state variable — the current budget balance — and a policy variable — water use q_{0t} — since the water trade decision would no longer be static. With this more complex approach we would see some 'precautionary saving' behaviour from the EWH (which we won't observe with our more pragmatic approach)

$$s_{it+1} = \min\{\max\{s_{it} - w_{it} - l_{it} + \lambda_i I_{t+1} + x_{it+1}, 0\}, K\}$$
$$s'_{it+1} = s_{it} - w_{it} - l_{it} + \lambda_i (I_{t+1} - Z_{t+1})$$
$$x_{it+1} = -Z_t \left(\frac{s'_{it+1}}{\sum_{i=1}^{n} s'_{it+1}} \right)$$
$$l_{it} = \left(\frac{s_{it}}{S_t} \right) L_t$$

Hughes (2015) shows that with out in-stream demands SWA tends to result in higher storage levels than CS and the planner outcome. On average, SWA achieves lower welfare than CS but the differences are often trivial.

6.6.3 Open access - OA

Here storage capacity is an open access resource. That is, there are no account limits and no evaporation loss deductions. Rather, all spills and losses are allocated in proportion to inflow shares (i.e., 'socialised'), such that user accounts follow

$$s_{it+1} = \min\{\max\{s_{it} - w_{it} + \lambda_i I_{t+1} + x_{it+1}, 0\}, K\}$$
$$x_{it+1} = \lambda_i (L_t + Z_t)$$

With out in-stream demands OA results in substantial over storage (higher storage levels than a planner's outcome). In most cases, OA achieves significantly lower welfare than CS and SWA (Hughes 2015).

6.6.4 No storage rights - NS

Here users have no storage rights. That is, any unused water is reallocated in proportion to inflow shares, so that user accounts follow

$$s_{it+1} = \lambda_i S_{t+1}$$

With no in-stream demands NS results in substantial under storage and is on average marginally outperformed by OA (Hughes 2015).

6.7 Solving the model

Given non-market interactions between users, the decentralised model is a stochastic game (i.e., a multi-agent dynamic programming problem).

Stochastic games present a number challenges: firstly defining a solution concept — such as a form of equilibrium — secondly establishing a method for computing solutions. Our approach is based around the ideas of 'multiple agent learning', where the economics of learning in games (Fudenberg and Levine 1998) meets the computer science methods of reinforcement learning (Sutton and Barto 1998).

Reinforcement learning (also known as approximate dynamic programming) is a subfield of machine learning, concerned with solving MDPs. Reinforcement learning algorithms optimise through simulation and so don't require an ex ante model of the 'environment' (i.e., probability transition and pay-off functions). Rather agent's 'learn' optimal policies by observing the outcomes — the payoffs and state transitions — of their actions.

Our approach is based on the method of 'Fitted Q iteration' (Ernst et al. 2005) a batch version of Q-learning. In fitted-Q-iteration a large number of state transition-action-payoff samples are accumulated through simulation, to which a Q or 'action-value' function is then fit. In a sense, the method translates a dynamic programming problem into a non-parametric regression (function approximation) problem. Here we use a version of tile coding (Sutton and Barto 1998).

In the multi-agent context this reinforcement learning update is performed iteratively, each time updating the policies of a random sample of users. The approach provides something of a middle ground between the rational expectations (i.e., dynamic programming) methods of modern macroeconomics and the simulation (i.e., genetic algorithm) approaches of agent based computational economics.

For detail on the solution methods see Hughes (2015).

7 Results

7.1 Central case

Social welfare results are presented in table 9, with profits in table 10 and environmental benefits in table 11[10]. Overall, CS-HL delivers the highest social welfare ($213.8m) and NS the lowest ($209.3m). All of the decentralised scenarios remain some distance from the planners solution at $217m.

In terms of storage levels (table 12) we find a similar outcome to the no in-stream flow case (Hughes 2015): OA leads to the highest storage levels, NS the lowest, while SWA to leads higher storage levels than CS. However, here all of the scenarios (with the exception of OA) lead to storage levels below the planner's solution.

This result can be explained by in-stream flow externalities: consumptive users do not take into account the environmental benefits of spills, such that mean storage levels are below the planner's solution. In contrast, OA — somewhat by accident — results in storage levels close to optimal (which explains why OA performs reasonably here well in terms of social welfare).

Table 9: Social welfare, ($m)

	Mean	SD	2.5th	25th	75th	97.5th
Planner	217.04	39.66	147.13	193.94	242.31	295.13
CS	211.95	41.95	129.05	185.08	239.25	290.80
SWA	212.55	41.87	124.26	188.52	238.29	292.46
OA	211.97	41.74	122.11	188.99	236.53	293.31
NS	209.32	46.75	103.24	180.49	241.28	290.60
CS-HL	213.79	41.35	129.43	187.66	241.45	288.90
SWA-HL	212.25	41.94	131.70	185.42	239.38	292.95

Table 10: Consumptive user profits $\sum_{i=1}^{n} u_{it}$ ($M)

	Mean	SD	2.5th	25th	75th	97.5th
Planner	183.34	19.77	143.03	175.36	195.97	208.53
CS	178.44	25.51	112.81	166.03	195.19	209.34
SWA	179.70	25.67	109.81	169.97	195.37	209.98
OA	177.80	25.59	104.25	170.41	193.39	206.02
NS	173.05	28.74	96.01	159.12	192.95	208.24
CS-HL	183.87	24.65	123.54	172.98	199.10	215.17
SWA-HL	179.19	25.16	115.56	167.80	195.68	209.11

[10]Throughout this section we focus on annual results

Table 11: Environmental benefits u_{0t} ($M)

	Mean	SD	2.5th	25th	75th	97.5th
Planner	33.70	30.59	1.27	7.54	52.38	107.42
CS	33.51	23.86	6.80	15.79	44.82	95.37
SWA	32.85	24.86	6.19	14.69	44.03	98.60
OA	34.17	26.93	6.10	14.76	45.42	107.98
NS	36.26	24.05	5.43	17.57	49.44	95.14
CS-HL	29.92	24.27	2.46	11.21	42.77	91.11
SWA-HL	33.06	24.27	7.39	15.02	44.33	97.25

Table 12: Storage, S_t (GL)

	Mean	SD	2.5th	25th	75th	97.5th
Planner	597.46	271.87	135.10	358.13	840.78	1,000.00
CS	543.91	245.71	138.82	349.56	732.56	1,000.00
SWA	591.43	260.22	159.14	372.07	818.91	1,000.00
OA	603.77	270.01	156.24	374.93	856.09	1,000.00
NS	479.91	247.28	119.87	279.79	648.85	1,000.00
CS-HL	521.77	252.47	119.08	320.66	712.38	1,000.00
SWA-HL	551.46	245.05	148.62	355.73	739.21	1,000.00

Table 13: Extraction, E_t (GL)

	Mean	SD	2.5th	25th	75th	97.5th
Planner	419.96	143.96	156.69	313.97	529.12	673.28
CS	403.00	162.54	115.60	276.86	527.69	710.95
SWA	402.83	142.32	109.93	305.04	510.50	631.05
OA	379.40	122.48	108.12	300.60	479.10	550.26
NS	380.74	178.52	88.21	237.36	517.84	696.86
CS-HL	439.85	182.26	114.93	296.61	595.00	759.24
SWA-HL	402.55	152.57	120.39	289.22	524.03	661.47

Table 14: Withdrawal, W_t (GL)

	Mean	SD	2.5th	25th	75th	97.5th
Planner	593.04	235.40	175.98	428.66	769.68	1,046.00
CS	623.62	276.31	184.64	402.25	827.29	1,159.04
SWA	594.04	216.52	180.00	433.92	776.20	940.58
OA	567.17	184.19	177.06	434.28	707.05	841.20
NS	633.27	295.48	161.80	389.02	871.71	1,156.83
CS-HL	622.55	270.71	173.67	401.73	849.10	1,083.04
SWA-HL	622.62	271.78	187.51	408.31	824.29	1,127.85

Table 15: Spills, Z_t (GL)

	Mean	SD	2.5th	25th	75th	97.5th
Planner	74.37	219.27	0.00	0.00	0.00	764.20
CS	47.05	190.04	0.00	0.00	0.00	593.39
SWA	73.04	235.08	0.00	0.00	0.00	836.49
OA	100.24	290.11	0.00	0.00	0.00	1,059.99
NS	40.95	178.19	0.00	0.00	0.00	559.64
CS-HL	48.87	194.61	0.00	0.00	0.00	643.88
SWA-HL	48.11	190.03	0.00	0.00	0.00	625.46

Figure 3: Mean profit versus mean environmental benefits, $\lambda_{0t} = 0.263$

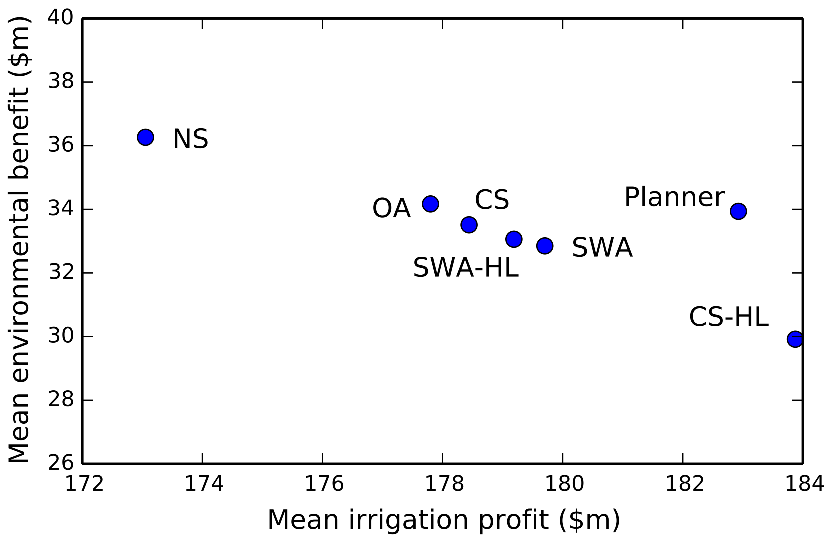

Figure 3 compares the scenarios in terms of mean profit and environmental benefits. Here we see a trade-off emerging between profits and the environment. The NS scenario leads to better environmental outcomes at the expense of profits, while CS-HL favours profits over environmental outcomes (for equal environmental shares, $\lambda_0 = 0.256$).

A clearer picture of the trade-off emerges when we vary the size of the environmental share. Here we solve the model with the central case parameters, but vary λ_0 over the range $[0.1, 0.5]$. This allows us to generate trade-off curves for each scenario (similar to those of Dudley et al. 1998), these curves are shown in figure 4. The mean social welfare results are summarised in table 16.

Figure 4: Profit-environmental benefit trade-off curves, for $\lambda_{0t} = 0.1$ to 0.5

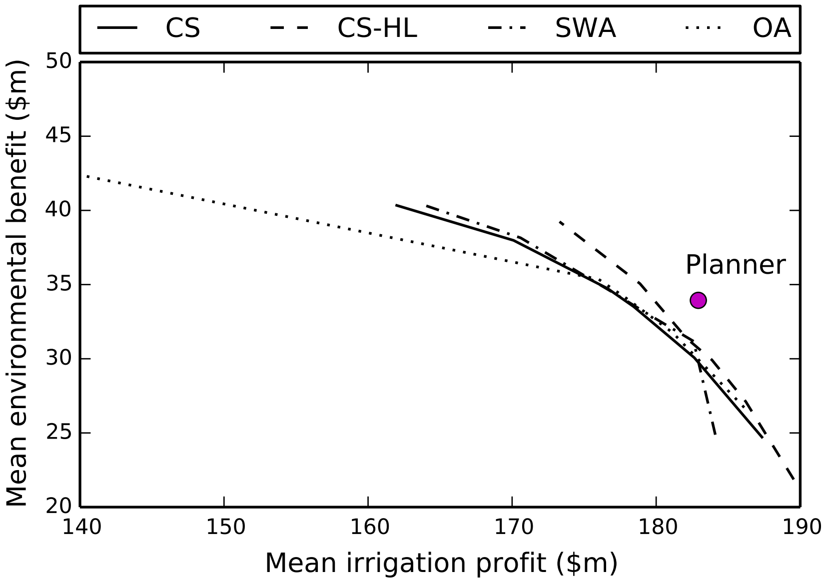

Table 16: Mean social welfare ($m) for $\lambda_{0t} \in [0.1, 0.2, 0.263, 0.3, 0.4, 0.5]$

	10	20	26.3	30	40	50
CS	212.06	212.72	211.95	211.37	208.08	202.33
SWA	212.09	213.78	212.55	211.48	208.74	204.13
OA	212.83	212.91	211.97	211.36	138.50	145.79
NS	209.56	210.46	209.32	207.94	203.58	190.39
CS-HL	211.42	213.33	213.79	213.50	213.94	212.53
SWA-HL	211.62	212.39	212.25	211.73	208.81	203.77

Now we see that CS-HL represents the 'frontier': the best of the decentralised scenarios. CS-HL with $\lambda_0 = 0.40$ is the best possible outcome yielding welfare of $213.9m with environmental benefits of $35.1m and profits of $176.5m. Since the EWH holds low reliability rights, a larger share is optimal. In contrast, under OA and NS lower (20 per cent) shares are optimal.

OA storage with $\lambda_0 \geq 0.4$ results in a dramatic reduction in welfare (table 16). With OA it becomes optimal for a large EWH to adopt a 'fill and spill' strategy: to make minimal withdrawals and accumulate storage reserves until the reservoir is full. With a full storage inflows spill uncontrolled downstream, leading to high environmental benefits but low profits: low reliability irrigation is essentially wiped out and high reliability users face frequent shortages. While such an extreme scenario is unlikely to occur in practice (see section 8), it casts doubts over the suitability of OA type storage rights for river systems with large EWHs.

7.1.1 Environmental water demand

As would be expected environmental water demand q_{0t} is more variable than consumptive demand (see table 17). Environmental demand is highest under the NS scenario, where the incentive is 'use it or lose it'. Under OA environmental demand is relatively low, however good river flows (and environmental benefits) are still achieved due to higher spills (table 15)

Table 17: Environmental use, q_{0t} (GL)

	Mean	SD	2.5th	25th	75th	97.5th
Planner	105.95	111.70	0.00	0.00	182.54	363.08
CS	124.27	112.36	0.00	29.87	197.40	378.71
SWA	100.61	83.15	0.00	27.69	162.53	276.73
OA	97.88	77.62	0.00	28.86	153.56	261.62
NS	149.85	113.02	0.96	50.03	232.60	372.60
CS-HL	94.20	86.59	0.00	16.21	167.75	270.75
SWA-HL	123.78	115.41	0.00	23.23	200.08	364.69

Figure 6 below shows mean environmental demand q_{0t} against storage S_t and inflow I_t for the CS, CS-HL, OA and SWA scenarios. Under OA and SWA we see lower EWH demand in high storage years — given the incentive to accumulate storage reserves and generate spills. Under CS we see demand for much larger environmental flows in wet years.

7.1.2 Environmental trade

Tables 18, 19 and 20 show the EWHs spot market trading patterns. In the long run, the EWH maintains an approximately balanced budget[11]. On average the EWH is a net seller of water in summer and net buyer in winter as would be expected.

[11]Given the approximate nature of the algorithm some small positive / negative balances are recorded

Figure 5: Mean environmental demand q_{0t} versus storage S_t and inflow I_t

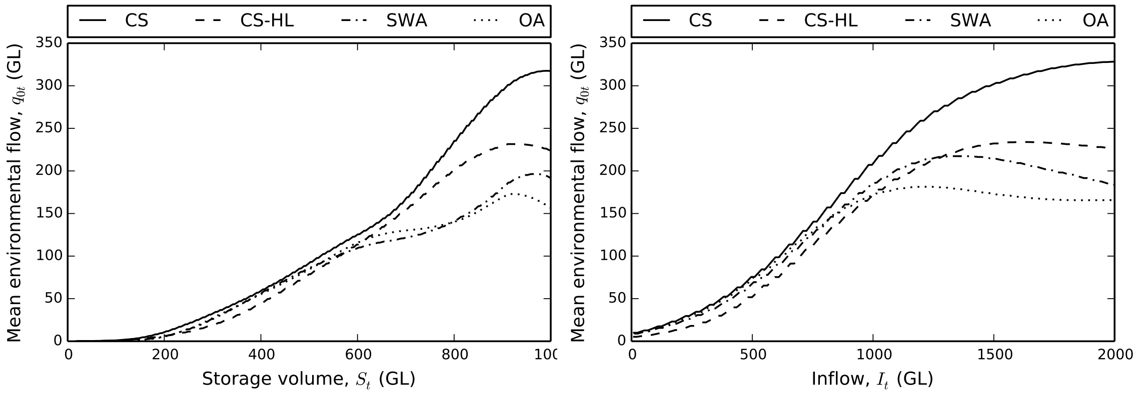

(a) versus storage

(b) versus inflow

However, trading patterns vary considerably across years: in some years the EWH is a net buyer and in others a net seller.

CS-HL leads to less EWH trade both across and within years: low priority rights are a good match for the EWH, minimising their trade requirements.

Table 18: Environmental trade - annual, $P_t(a_{0t} - q_{0t})$ (\$m)

	Mean	SD	2.5th	25th	75th	97.5th
Planner	0.00	0.00	0.00	0.00	0.00	0.00
CS	-0.08	4.83	-11.34	-2.52	3.31	7.91
SWA	-0.05	5.09	-11.70	-2.55	3.21	8.85
OA	-0.11	5.44	-12.02	-2.86	2.89	10.24
NS	0.11	4.14	-9.00	-2.10	2.72	7.52
CS-HL	0.05	1.93	-5.26	0.00	0.80	3.49
SWA-HL	-0.11	5.33	-12.08	-3.02	3.71	8.56

Table 19: Environmental trade - summer, $P_t(a_{0t} - q_{0t})$ (\$m)

	Mean	SD	2.5th	25th	75th	97.5th
Planner	0.00	0.00	0.00	0.00	0.00	0.00
CS	2.09	3.10	-3.43	0.00	4.50	8.22
SWA	1.75	3.91	-6.59	0.00	4.41	9.35
OA	2.20	3.91	-5.40	0.00	4.31	10.59
NS	2.13	3.19	-3.57	0.00	4.59	8.03
CS-HL	0.30	1.76	-4.58	0.00	0.98	3.61
SWA-HL	2.25	3.47	-3.84	0.00	4.98	8.80

Figure 6 below shows mean environmental trade value $P_t(q_{0t} - a_{0t})$ against storage S_t and inflow I_t. Here we essentially see the 'counter cyclical' type trading pattern observed in previous studies (see Kirby et al. 2006), where the EWH is selling water during 'dry' periods and buying in 'wet'.

Table 20: Environmental trade - winter, $P_t(a_{0t} - q_{0t})$ ($m)

	Mean	SD	2.5th	25th	75th	97.5th
Planner	0.00	0.00	0.00	0.00	0.00	0.00
CS	-2.17	3.19	-11.17	-3.32	0.00	0.00
SWA	-1.80	2.85	-9.92	-2.73	0.00	0.00
OA	-2.31	3.31	-11.48	-3.72	0.00	0.00
NS	-2.02	2.61	-9.08	-3.29	0.00	0.00
CS-HL	-0.25	0.84	-3.08	0.00	0.00	0.00
SWA-HL	-2.36	3.45	-11.91	-3.76	0.00	0.00

However, in very high inflow years the EWH is less likely to buy (as it can rely on spills). Similarly, in years with very low storage the EWH is less likely to sell water. CS-HL results in a less trade on average and a different pattern of trade, with the EWH on average a marginal net buyer in low storage years

Figure 6: Mean environmental trade $P_t(q_{0t} - a_{0t})$ versus storage S_t and inflow I_t

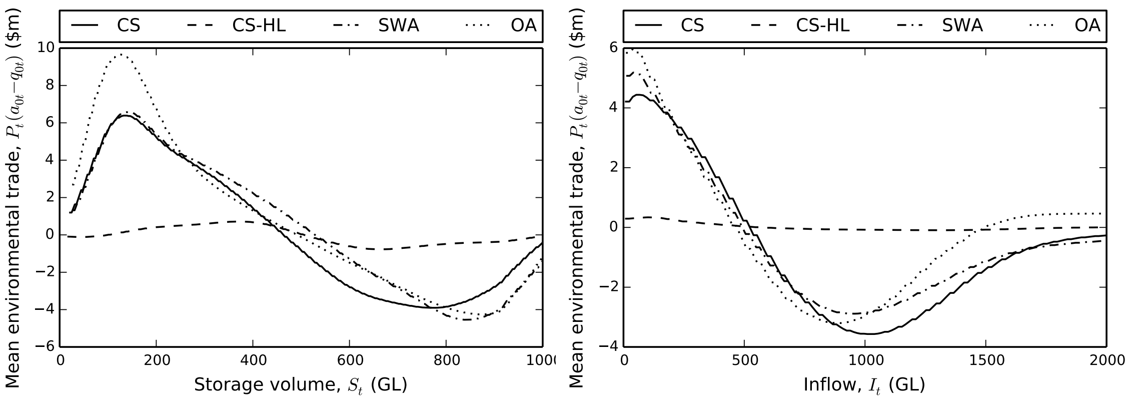

(a) versus storage

(b) versus inflow

Hughes (2015) presents the results of a no-trade scenario and computes the gains from spot market trade (figure 7). With environmental demands the gains from spot market trade are significant (in the order of $6m a year) even in scenarios with well defined storage rights and/or priority rights.

Figure 7: Gains from trade ($M)

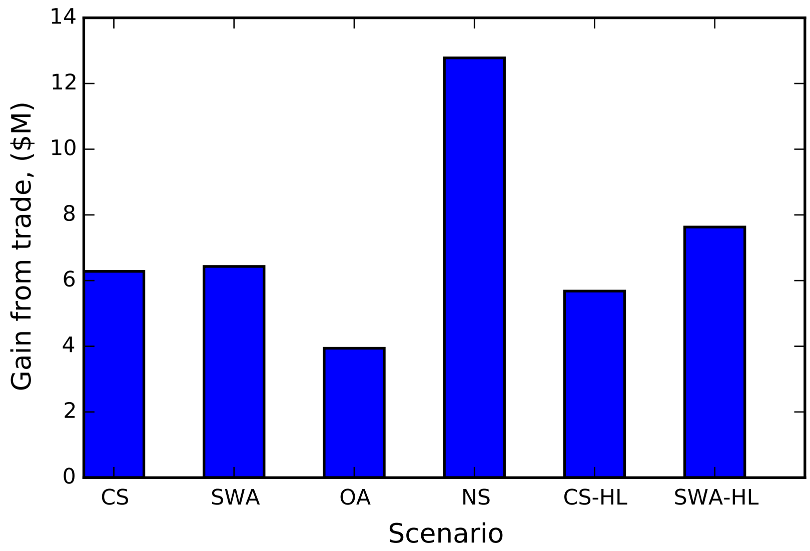

7.2 General case

Here we draw 550 parameter sets and solve the model for the CS, CS-HL, SWA and OA scenarios. Below we summarise the results for mean social welfare, profit, environmental benefits and storage volumes, we also present indexes relative to the CS scenario.

7.2.1 Social welfare

Mean social welfare results are summarised in tables 21 and 22 and figure 8. Here, CS-HL is the most frequently preferred scenario (198 of 547 complete runs), followed by OA (132), SWA (120) and CS (97).

Figure 8: Social welfare index, general case

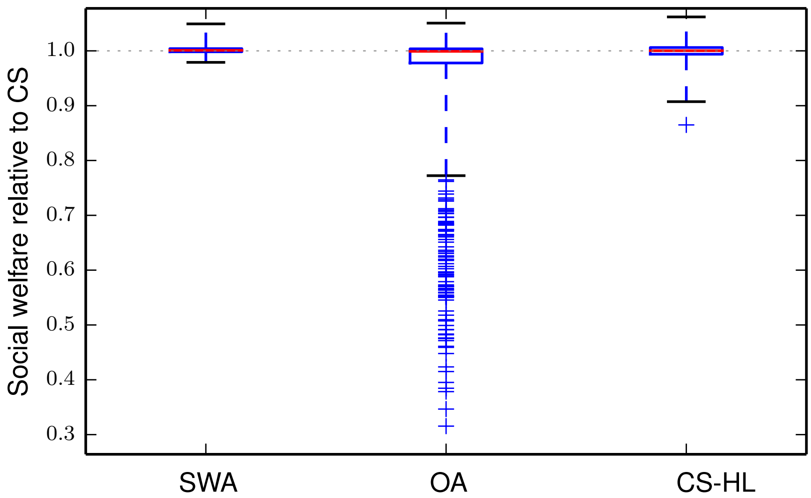

Table 21: Mean social welfare ($m), general case

	Mean	Min	Q1	Q3	Max
CS	216.85	59.01	151.85	280.52	430.41
SWA	217.16	58.96	151.85	280.96	431.93
OA	195.18	50.80	135.42	244.90	431.19
CS-HL	216.82	59.92	151.89	280.15	432.02

Table 22: Social welfare index, general case

	Mean	Min	Q1	Q3	Max
CS	1.00	1.00	1.00	1.00	1.00
SWA	1.00	0.98	1.00	1.00	1.05
OA	0.92	0.25	0.98	1.00	1.05
CS-HL	1.00	0.86	0.99	1.01	1.06

As in the central case, OA can lead to extremely low welfare under certain conditions. Figure 9 plots OA welfare relative to CS, against the environmental share λ_0 and mean inflow relative to capacity $E[I_t]/K$. OA tends to perform very poorly where inflow is high relative to capacity (spills are frequent) and the environmental share is large.

Figure 9: OA welfare index, against $E[I_t]/K$ and λ_0

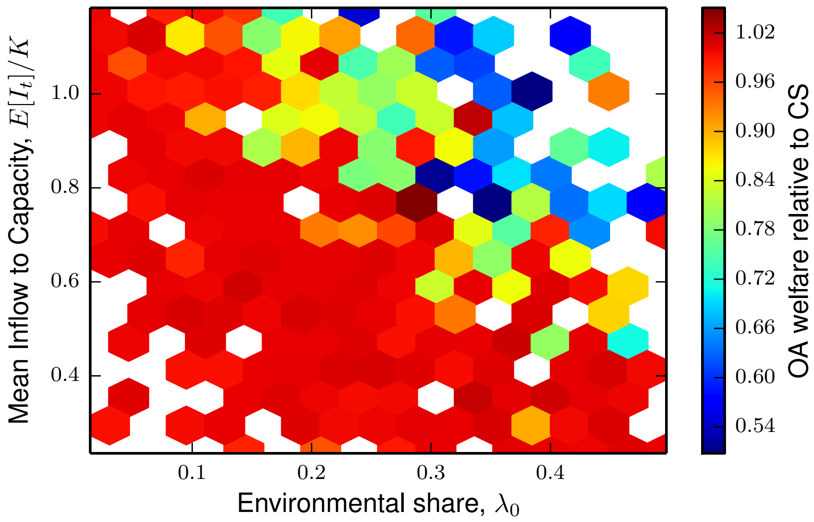

Next we regress our mean welfare index against the model parameters. The most important parameters (table 23) are the environmental value ($b_\$/\bar{I}$), mean inflow ($E[I_t]/K$), the amount of high reliability demand (n_{high}) and the environmental inflow share 'shock' ($\lambda_0 - \hat{\lambda}_0$).

The effect of these parameters is summarised in figure 10. Together, $b_\$/\bar{I}$ and $\lambda_0 - \hat{\lambda}_0$ determine the environmental share λ_0. Figure 10 shows how the performance of OA quickly deteriorates for scenarios with high λ_0 (high $b_\$/\bar{I}$ and $\lambda_0 - \hat{\lambda}_0$). OA also performs poorly for high values of $E[I_t]/K$ or n_{high}.

On average, CS-HL outperforms alternatives when $\lambda_0 > \hat{\lambda}_0$. As we found in the central case, a CS-HL scenario in which the environment holds a large share of low reliability rights appears to be ideal[12].

[12]Our assumption $\lambda_0 \sim N(\hat{\lambda}_0, 0.05)$ biases our general case results against CS-HL, since the optimal λ_0 for CS-HL will be greater than $\hat{\lambda}_0$. If we were to compute optimal inflow shares (as we did in the central case trade-off results) CS-HL would be more frequently prefered than it is here.

Table 23: Social welfare index regression, parameter importance

	Importance
$\frac{b_{\$}}{\bar{I}}$	12.90
$E[I]/K$	12.52
$\lambda_0 - \hat{\lambda}_0$	6.52
$\frac{\mathscr{A}_{low}}{E[I]/K}$	4.65
n_{high}	4.60
σ_{e0}	4.58
c_v	4.09
δ_a	3.73
ρ_I	3.69
δ_R	3.66
b_1	3.55
μ_ω	3.48
σ_η	3.41
δ_{Eb}	3.24
α	3.19
τ	3.09
δ_{Ea}	2.88
$\Lambda_{high} - \hat{\Lambda}_{high}$	2.84
ρ_e	2.83
$\Lambda_{high}^{CS-HL} - \hat{\Lambda}_{high}^{CS-HL}$	2.79
σ_ω	2.61
ω_δ	2.58
$\delta 0$	2.55

Figure 10: Social welfare index regression results

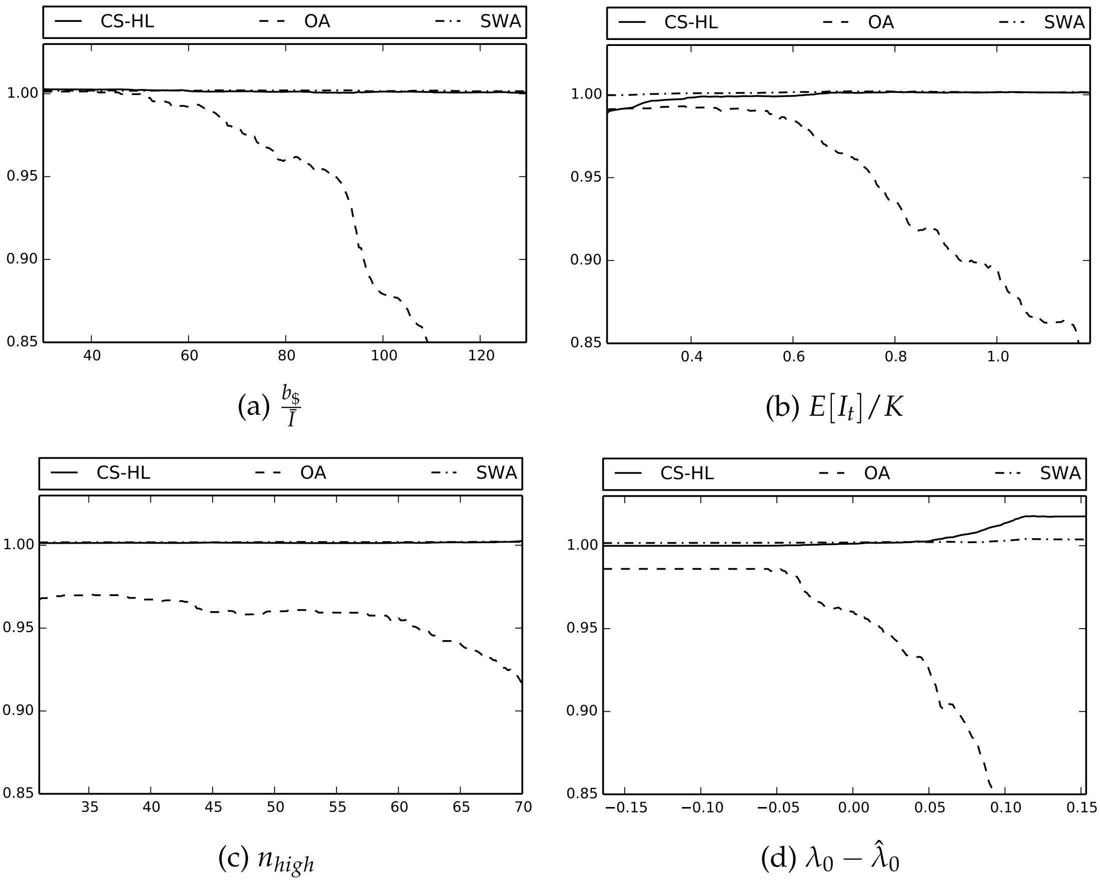

(a) $\frac{b_\$}{\bar{I}}$

(b) $E[I_t]/K$

(c) n_{high}

(d) $\lambda_0 - \hat{\lambda}_0$

7.2.2 Profit

Mean profits are summarised in tables 24 and 25 and figure 11. As in the central case, CS-HL tends to result in higher profits and OA in lower profits for equivalent λ_0. Figure 11 shows the long tail of low profit outcomes under OA. As in the central case, these low welfare outcomes are a result of the EWH adopting a 'fill and spill' strategy leading to low profits but high environmental benefits (figure 12), storage levels (figure 13) and spills.

Figure 11: Mean profit index, general case

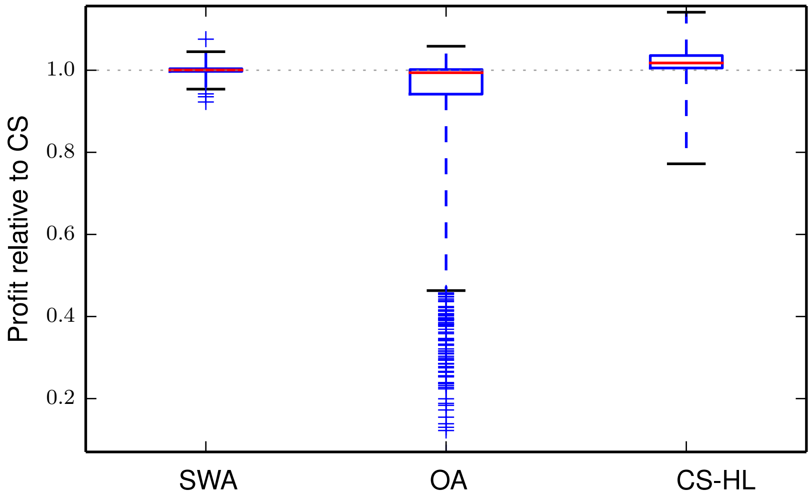

Table 24: Mean profit ($m), general case

	Mean	Min	Q1	Q3	Max
CS	183.16	51.18	129.77	238.25	379.04
SWA	183.10	50.64	129.22	238.59	378.90
OA	154.93	8.35	91.14	210.61	379.75
CS-HL	186.20	53.89	133.32	240.88	381.40

Table 25: Mean profit index, general case

	Mean	Min	Q1	Q3	Max
CS	1.00	1.00	1.00	1.00	1.00
SWA	1.00	0.92	1.00	1.00	1.08
OA	0.87	0.05	0.94	1.00	1.06
CS-HL	1.02	0.77	1.01	1.04	1.14

7.2.3 Environmental benefits

Mean environmental benefits are summarised in tables 26 and 27 and figure 12. Again, we see that OA favours the environment and CS-HL the consumptive users, given equal EWH shares. On average, SWA generates slightly higher environmental benefits than CS.

Figure 12: Mean environmental benefit, general case

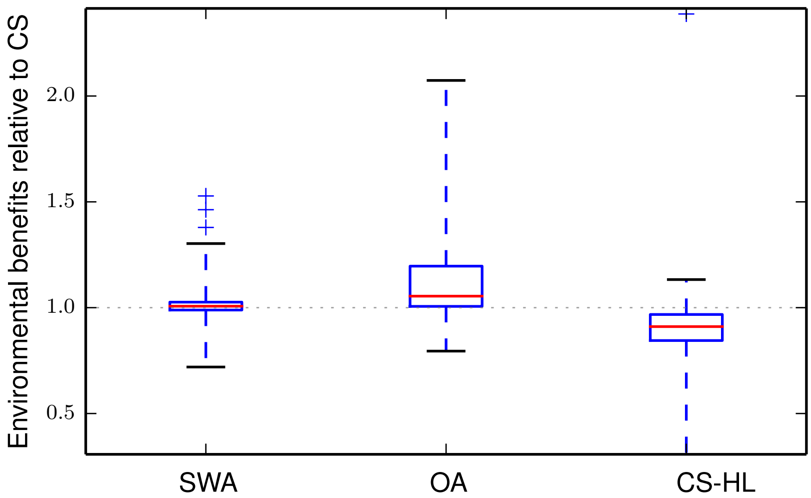

Table 26: Mean environmental benefits ($m), general case

	Mean	Min	Q1	Q3	Max
CS	33.69	3.21	17.10	44.89	102.23
SWA	34.06	3.27	17.27	45.79	99.38
OA	40.25	3.36	17.94	55.05	134.74
CS-HL	30.62	1.98	14.99	41.40	95.02

Table 27: Mean environmental benefit index, general case

	Mean	Min	Q1	Q3	Max
CS	1.00	1.00	1.00	1.00	1.00
SWA	1.01	0.72	0.99	1.03	1.53
OA	1.14	0.79	1.01	1.20	2.07
CS-HL	0.89	0.30	0.84	0.97	2.39

7.2.4 Storage

Mean storage levels are summarised in tables 28 and 29 and figure 13. On average OA leads to storage levels 17 per cent higher than CS, SWA 3 per cent higher and CS-HL 2 per cent lower.

Figure 13: Mean storage index, general case

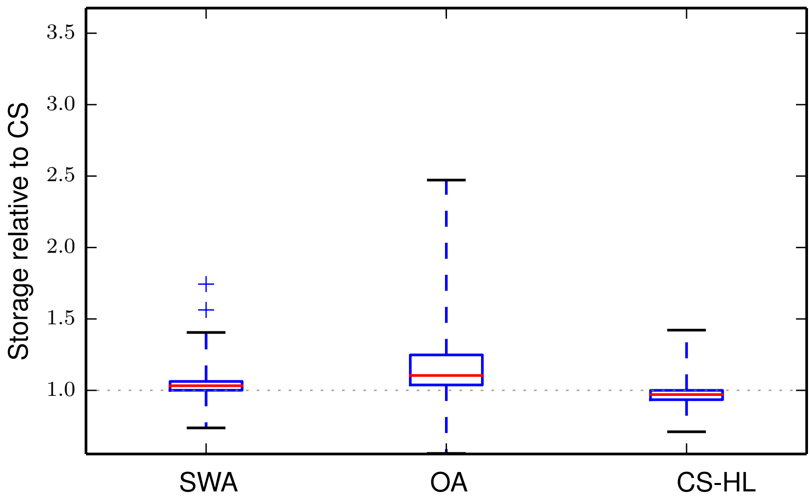

Table 28: Mean storage (GL), general case

	Mean	Min	Q1	Q3	Max
CS	554.19	177.84	479.50	652.37	864.64
SWA	572.90	176.79	490.79	678.79	990.96
OA	657.46	168.36	520.11	815.31	999.92
CS-HL	536.70	204.14	459.48	625.67	964.27

Table 29: Mean storage index, general case

	Mean	Min	Q1	Q3	Max
CS	1.00	1.00	1.00	1.00	1.00
SWA	1.03	0.74	1.00	1.06	1.74
OA	1.18	0.56	1.04	1.25	3.68
CS-HL	0.97	0.71	0.93	1.00	1.42

8 Conclusions

8.1 Storage rights

With environmental demands, storage rights generally result in lower storage levels than a planner's solution: given most users ignore the environmental benefits of spills. Among the storage right options, we observed similar relative effects to Hughes (2015): OA leads to the highest storage, NS the lowest and SWA to slightly higher storage than CS.

The nature of the welfare effects changes significantly. In some cases OA performs well on account of higher spills and environmental benefits. In other cases — particularly where the EWH's share of water rights is large and inflows are high relative to capacity — it can lead to disastrously bad outcomes.

Under OA, it can be optimal for large EWHs to adopt a 'fill and spill' strategy, where they deliberately accumulate storage reserves to generate spills which benefit the environment, but limit the consumptive supply of water.

The most extreme outcomes — where irrigation is essentially wiped out — are unlikely to occur in the MDB for at least three reasons. Firstly, there are no rivers with pure OA storage rights. Second, the CEWH has a smaller share than is required to generate this outcome from the model (30 to 40 per cent). Finally, such extreme behaviour on the part of the EWH would not be politically feasible.

Regardless, the results are enough to recommend against the adoption of OA storage rights in the presence of EWHs. In contrast, CS is a robust property rights system: it is the most frequently preferred option and it performs well in almost all types of river systems, both with and without EWHs.

8.2 Inflow rights

Hughes (2015) showed how storage rights help mitigate trade requirements. With well defined storage rights (i.e., CS) the gains from trade are small and the benefits of priority rights are negligible (and in some cases negative).

The introduction of an EHW changes this result. Here the gains from spot market trade are large. The EWH trading patterns in our model are more or less consistent with existing studies. EWH trading is frequently 'counter cyclical': the EWH sells water in dry periods and buys during wet.

In this context, priority rights are found to offer a tangible improvement over proportional rights. Low priority rights are a good match for the demands of EWHs and significantly reduce their trade requirements.

8.3 The ideal rights system

In the introduction we asked: which form of water property rights is ideal in the presence of a large EWH? The short answer is CS-HL: capacity sharing with priority inflow rights. A CS-HL scenario in which the EWH holds a larger share (40 per cent) of low priority rights was the ideal outcome in the central case. In the general case, CS-HL was the most frequently preferred scenario.

At a very high level, CS-HL is the property rights system on which most parts of the MDB appear to be converging (see Hughes et al. 2013).

8.4 Future research

In this study, all of the decentralised scenarios remain some distance from an optimal planner's outcome. This raises the question of whether a rules based system or more likely a mix of rules and decentralisation could outperform a pure market approach.

One area for future research, would be testing a combination of environmental flow rules and a discretionary EWH. Another would be including flood mitigation objectives and rules.

*References

Ashton, D. and Oliver, M. 2012, *An economic survey of irrigation farms in the Murray-Darling Basin: industry overview and region profiles, 2009-10*, Australian Bureau of Agricultural and Resource Economics and Sciences.

Australian National Committee on Large Dams 2013, 'Register of large dams in Australia'.
<www.ancold.org.au>

Beare, S. C., Hinde, R., Heaney, A., Che, N. and Hillman, T. 2006, Meeting environmental water demands under uncertainty, *in* 'Presented at the 3rd World Congress of Environmental and Resource Economists'.

Brennan, D. 2011, 'Economic perspectives of the proposed basin plan for the southern connected murray'.
<http://press.anu.edu.au//wp-content/uploads/2011/08/whole5.pdf>

Bureau of Meterology 2013, 'Climate data online'.
<www.bom.gov.au>

Dudley, N., Arthington, A., Scott, B. and van der Lee, J. 1998, *Integrating environmental and irrigation water allocation under uncertainty*, Centre for Water Policy Reasearch.

Dudley, N. J. 1988, 'A single decision-maker approach to irrigation reservoir and farm management decision making', *Water Resources Research*.

Dudley, N. J. and Musgrave, W. F. 1988, 'Capacity sharing of water reservoirs', *Water Resources Research*, vol. 24, no. 5, pp. 649–658.

Ernst, D., Geurts, P. and Wehenkel, L. 2005, 'Tree-based batch mode reinforcement learning', , vol. pp. 503–556.

Fudenberg, D. and Levine, D. K. 1998, *The theory of learning in games*, Vol. 2, MIT press.

Grafton, R. Q., Chu, H. L., Stewardson, M. and Kompas, T. 2011, 'Optimal dynamic water allocation: Irrigation extractions and environmental tradeoffs in the Murray River, Australia', *Water Resources Research*.

Heaney, A., Beare, S. and Brennan, D. 2011, Managing Environmental Flow Objectives Under Uncertainty, *in* 'Annual Conference of the Australian Agricultural and Research Economics Society'.

Hughes, N. 2015, Water property rights in regulated rivers, PhD thesis, Australian National University.

Hughes, N., Gibbs, C., Dahl, A., Tregeagle, D., Sanders, O. and Goesch, T. 2013, *Storage rights and water allocation arrangements in the Murray-Darling Basin*, Australian Bureau of Agricultural and Resource Economics and Sciences.

Kirby, M., Qureshi, E., Mohammed, M. and Dyack, B. 2006, 'CATCHMENT BEHAVIOR AND COUNTER-CYCLICAL WATER TRADE: AN INTEGRATED MODEL', *Natural Resource Modeling*, vol. 19, no. 4, pp. 483–510.

Lund, J. R. 2006, 'Drought storage allocation rules for surface reservoir systems', *Journal of water resources planning and management*, vol. 132, no. 5, pp. 395–397.

National Water Comission 2011, *Australian water markets report 2010-11*.

Sutton, R. and Barto, A. 1998, *Reinforcement learning: An introduction*, Vol. 1, Cambridge University Press.

Truong, C. 2011, 'An Analysis of Storage Capacity Reallocation Impacts on the Irrigation Sector', *Environmental and Resource Economics*, vol. 51, no. 1, pp. 1–19.