Water storage rights:
Decentralising reservoir operation

Neal Hughes

Australian National University

August 12, 2014

Abstract

Given incomplete information it can be difficult for central planners to design optimal reservoir storage policies. An increasingly popular option, is to decentralise reservoir operation by defining storage property rights. One approach is 'capacity sharing' (Dudley and Musgrave 1988) where water users are assigned shares in the storage capacity and inflow of a reservoir.

As with spatial water markets, decentralised storage is subject to externality problems. User storage decisions can have external effects on other users, via their effect on spills and evaporation losses. Even sophisticated approaches like capacity sharing are subject to externalities, including the problem of 'internal spills'.

This paper presents a model of a regulated river system, in which a large number of water users (e.g., irrigation farmers) make their own storage decisions as well as engaging in a water spot market. The purpose of the model is to evaluate alternative approaches to storage property rights. Given the presence of externalities our model represents a stochastic game. We solve this model numerically using reinforcement learning methods.

We find that capacity sharing achieves close to optimal levels of welfare and is the preferred approach for almost all considered parameterisations of the model. Poorly specified storage rights impose welfare losses by inducing storage behaviour that departs significantly from the optimum. In particular, open access storage — where spills and evaporation losses are socialised — can lead to dramatic over storage.

1 Introduction

Storage rights allow water users to hold private storage reserves in public reservoirs. As such, storage rights at least partially decentralise reservoir operation. User level storage rights are now common place in the Australian Murray-Darling Basin (MDB). Similar arrangements exist in a number of western US rivers. Recently a form of storage right has even emerged on the Colorado river (Hughes 2013).

Surface water storage rights have been examined in a number of Australian studies (Dudley and Musgrave 1988, Brennan 2008, Hughes and Goesch 2009*b*). Existing research emphasises the limitations of central control of reservoirs. The idea being that storage rights may improve the inter-temporal allocation of water just as trading may improve the spatial allocation.

The complexity of water makes defining storage rights difficult. Storage capacity represents a 'congestible good' (Randall 1983), switching from non-rival to rival as storages fill. Further, storage losses vary non-linearly with volumes. As a result of these complexities, storage rights are never completely exclusive.

In this paper, we compare a number of approaches to water storage rights observed in practice, including the capacity sharing model advocated by Dudley (1988*a*). These alternatives differ on two dimensions: how they reflect the storage capacity constraint (spills) and how they reflect evaporation losses. Such a comparison has not yet been attempted in the literature, primarily because it requires a decentralised model.

In this paper, we present a decentralised (multi-agent) model of a regulated river system. In this model each user makes forward looking storage decisions, while also engaging in a water spot market. Formally the model is a stochastic game: each user is faced with a Markov decision process (MDP) where the payoffs and state transition probabilities are dependent of the actions of other users.

We solve this model numerically with a relatively novel application of reinforcement learning. Reinforcement learning is a sub-field of machine learning, which provides a range of algorithms for solving MDPs by simulation. The model is solved for a large number of parameterisations, using parameter distributions reflective of the Australian MDB.

The goal of this chapter is to address the following questions: which system of storage rights maximises social welfare? How do the systems affect user storage behaviour and therefore aggregate storage volumes? How do the systems affect the distribution of welfare? And finally, how do the answers to these questions vary with the parameterisation of the model?

2 The problem

Below we introduce the general set up of the model, in Section ?? we present a parametric version.

This paper is concerned with an abstract regulated river system summarised in Figure 1. This system involves a single reservoir which receives stochastic inflows and delivers water to n consumptive users located at a single demand node.

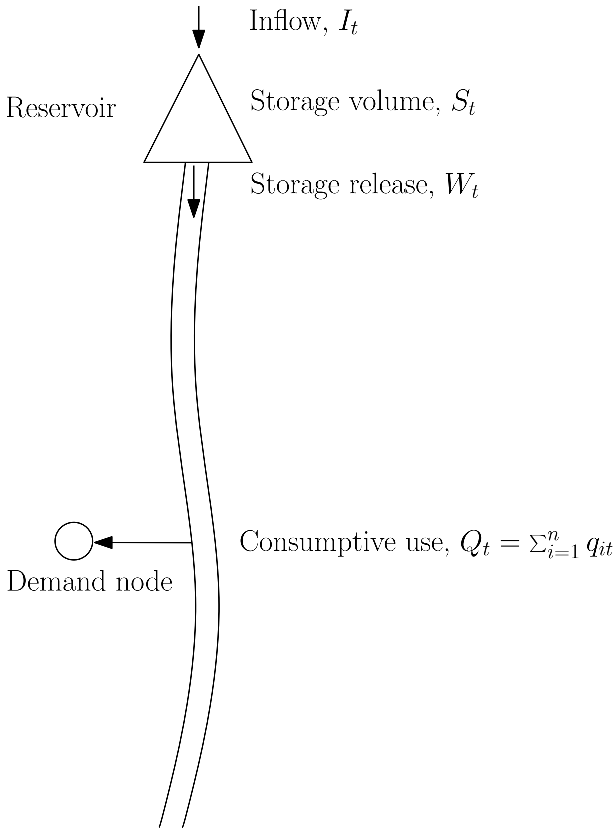

Figure 1: A simple regulated river system

The model is in discrete time with an infinite time horizon. The storage has transition rule

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

$$0 \leq W_t \leq S_t$$

Here S_t is the storage volume, I_t the stochastic inflow, W_t the storage release (withdrawal), L_t the storage evaporation loss and K the fixed storage capacity. Storage losses are a concave function $\mathcal{L}_0$ of storage contents.

$$L_t = \mathcal{L}_0(S_t) \in [0, S_t]$$

For convenience we define storage spills Z_{t+1} as

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

q_{it} is the volume of water consumed by user i in period t. Total consumptive water use $Q_t = \sum_{i=1}^{n} q_{it}$ is constrained by the volume delivered to the demand node

$$Q_t = \sum_{i=1}^{n} q_{it}$$

$$Q_t \leq W_t - f_1(W_t)$$

where $\mathcal{L}_1$ is a loss function, $\mathcal{L}_1(W_t) \in [0, W_t]$.

The users have payoff functions $\pi_i(q_{it}, I_t)$. Here the inflow I_t acts as a proxy for moisture availability (e.g., irrigation area rainfall). π_i is concave in q_{it} and I_t. q_{it} and I_t are substitutes.

2.1 The planners problem

The planners problem is to set policy variables W_t and q_{it}, conditional on state variables S_t, I_t and e_{it}, so as to maximise the expected discounted sum of user payoffs

$$\max_{\{q_{it}, W_t\}_{t=0}^{t=\infty}} E\left\{\sum_{t=0}^{\infty} \beta^t \sum_{i=1}^{n} \pi_i(q_{it}, I_t)\right\}$$

subject to to the above water supply constraints.

2.2 The decentralised problem

In the decentralised version property rights are defined, which facilitate both a water spot market and user level storage decisions. Below we outline the general form of the problem, before detailing specific scenarios in Section 3.

2.2.1 The property rights framework

Here each user controls their own 'water account'. Each period these accounts are credited with a share λ_i of inflow and debited for user withdrawals w_{it}. The evolution of user account balances s_{it} follows

$$s_{it+1} = \min\{\max\{s_{it} - w_{it} - l_{it} + \lambda_i I_{t+1} + x_{it+1}, 0\}, k_{it}\}$$

$$w_{it} \leq s_{it}$$

$$\sum_{i=1}^{n} \lambda_i = 1, \quad \sum_{i=1}^{n} s_{it} = S_t, \quad \sum_{i=1}^{n} w_{it} = W_t, \quad \sum_{i=1}^{n} l_{it} = L_t$$

where l_{it} are user storage loss deductions, k_{it} are user account limits and x_{it} are the '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.

A storage rights system is defined by the specification of l_{it}, k_{it} and x_{it}. A number of approaches are defined in the following section. For now, note that x_{it} can be a rather complicated function of the storage balances and withdrawals of all users $\boldsymbol{s_t} = (s_{1t}, s_{2t}, ..., s_{nt})$ and $\boldsymbol{w_t} = (w_{1t}, w_{2t}, ..., w_{nt})$ as well as physical quantities S_t, W_t, L_t and I_t.

2.2.2 The water spot market

Users receive water allocations a_{it} adjusted for delivery losses. For now we assume linear delivery losses

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

Water allocations can be used or traded on the spot market, subject to the market clearing condition

$$\sum_{i=1}^{n} q_{it} = \sum_{i=1}^{n} a_{it}$$

The spot market is subject to a positive transfer cost τ, such that user payoffs u_{it} are defined

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

where P_t is the market price for water. The effective user demand functions $\tilde{d}_h$ can then be defined as

$$q_{it} = \tilde{d}_h(P_t, \tau, a_{it}, I_t, e_{it}) = \begin{cases} d_h(P_t, I_t, e_{it}) & \text{if } p_{it} \leq P_{it} \\ a_{it} & \text{if } P_t < p_{it} < P_{it} + \tau \\ d_h(P_t + \tau, I_t, e_{it}) & \text{if } p_{it} \geq P_{it} + \tau \end{cases}$$

$$p_{it} = d_h^{-1}(a_{it}, I_t, e_{it})$$

$$d_h^{-1}(q_{it}, I_t) = \max\left\{\frac{\partial \pi_i}{\partial q_{it}},\ 0\right\}$$

2.2.3 The users' problem

The users' problem is to determine w_{it} and q_{it} each period in order to maximise their expected discounted payoff

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

subject to the above water accounting constraints, the behaviour of the other agents and the physical constraints as detailed in the planners problem. Formally this problem is a stochastic game, solution concepts and methods are detailed later in Section 7.

3 Policy scenarios

Our goal is to use the above framework to compare a number of water storage right systems. These systems differ on two dimensions: how they reflect the storage capacity constraint (how k_{it} and x_{it} are defined) and how they reflect evaporation losses (how l_{it} is defined).

In Section 5 we provide mode detail on how these scenarios relate to the storage rights systems of the western US and the Australian MDB.

3.1 Storage capacity rights (capacity sharing) - CS

Here each user is assigned both a share of total storage capacity and a share of inflow. For now we assume storage and inflow shares are equal, so that users' accounts follow

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

This scenario is representative of capacity sharing as proposed by Dudley and Musgrave (1988). Capacity sharing has been implemented at the irrigator level in two Queensland MDB irrigation areas (Hughes and Goesch 2009*b*). The state level property rights of NSW and VIC on the Murray river are another example of this approach. Similar storage right arrangements exist in Northern NSW (Hughes et al. 2013) and in some US systems, specifically the Texas Lower Rio-Grande (see Section 5 for more detail).

While capacity sharing lessens storage externalities in comparison with alternatives it does not eliminate them. With capacity sharing we have 'internal spills': where a users account reaches its limit and excess inflow is forfeited and reallocated to other users.

Here x_{it+1} represents user $i's$ share of the total pool of internal spills at time t. $x_{it+1} = 0$ for all i if no accounts reach capacity or if all accounts reach capacity (and the storage physically spills) . With only two users $\{i, j\}$, x_{it+1} is defined:

$$x_{it+1} = \begin{cases} \max\{\lambda_j.I_{t+1} - (\lambda_j K - s_{jt} + w_{jt} + l_{jt}), 0\} & \text{if } Z_{t+1} = 0 \\ 0 & \text{otherwise} \end{cases}$$

With a large number of users calculating x_{it+1} is complicated since an initial reallocation of internal spills may fill further accounts creating more internal spills and so on. In this case x_{it+1} can be calculated iteratively.

3.2 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 spill deductions in proportion to their account volumes, specifically:

$$s_{it+1} = \min\{\max\{s_{it} - w_{it} - l_{it} + \lambda_i I_{t+1} + x_{it+1}, 0\}, K\}$$
$$x_{it+1} = -Z_t \left(\frac{s_{it} - w_{it} - l_{it} + \lambda_i (I_{t+1} - Z_{t+1})}{K} \right)$$

Such an approach is adopted in Northern Victoria under the banner of 'spillable water accounts'. Similar approaches also exist in the US (see Section 5.

3.3 Open access storage (unlimited carryover) - OA

Here storage capacity is treated as an open access resource. Specifically, there are no limits on users storage volumes and no loss deductions. Rather all spills and losses are allocated based on 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 a large number of users, open access will result in too much water being stored. While there are few examples of pure open access storage in the MDB, many systems can approach open access under certain conditions, if rules designed to limit storage fail to bind (Hughes et al. 2013).

3.4 No storage access (use it or lose it) - NS

In this scenario users have no storage rights. Specifically any water unused at the end of the period is reallocated across all users in proportion to inflow shares, so that user accounts are simply:

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

With a large number of users the incentive is to 'use it or lose it'. That is to consume or trade all water allocations or at least until the users marginal value for water or the effective market price is zero.

This scenario is broadly reflective of the MDB prior to the introduction of storage rights. While central storage policies were in place, typically the maximum release volume was well in excess of demand, such that significantly more water was allocated than was ever used (MDBMC 1995).

3.5 Storage loss deduction - LD

For the capacity sharing (CS) and spillable water account (SWA) scenarios we consider two approaches to storage loss deductions. The first is to allocate total losses in proportion to user account balances:

$$l_{it} = \left(\frac{s_{it}}{S_t}\right) L_t$$

This is essentially the approach adopted to storage deductions in southern Queensland (Hughes and Goesch 2009*b*). A similar approach is adopted in northern Victoria.

3.6 Socialised storage loss - SL

The second and more common approach (Hughes and Goesch 2009*b*) is to socialise storage losses, that is to allocate total losses in proportion to user inflow shares

$$l_{it} = \lambda_i L_t$$

4 Literature

Existing literature on surface water storage relies heavily on social planner models. While storage rights have been considered in a number of Australian studies (Dudley and Musgrave 1988, Brennan 2008, Hughes and Goesch 2009*b*), they have rarely been explicitly modelled. Instead, these authors generally argue that optimal scenarios are broadly reflective of a storage rights outcome.

The literature on surface water storage rights begins with the work of Norman Dudley (Dudley 1988*b*;*a*; 1992; 1999), a long time advocate for capacity sharing:

> [Capacity sharing] is a property rights structure and institutional arrangement that allows multiple water users to each act as if they had their own small reservoir on their own small stream. It does so by providing each user, or small group of users, of reservoir water with long-term rights to a percentage of reservoir inflows and percentage of empty reservoir capacity or space in which to store those inflows, and from which to control releases. Their reservoir releases through time can be managed according to their particular supply reliability preferences....Their probabilities of water supply from their streamflow shares can be calculated directly from historical or synthesised streamflow data. (Dudley 1999; pp. 243)

Here Dudley and Musgrave (1988) identify two advantages of capacity sharing: closer alignment of storage policy with user preferences and a reduction in policy uncertainty. However, Dudley and Musgrave (1988) are careful to acknowledge that under capacity sharing, users are not entirely independent:

> [Capacity sharing users] are like a bank depositor who cannot incur a negative balance, cannot accumulate deposits in excess of a maximum and cannot control amount or timing of deposits. Instead, deposits are made according to a stochastic process.... However, beyond these stochastic deposits... there may be extra deposits made periodically

> to a depositor's account because of the heterogeneous behavior of all depositors. (Dudley and Musgrave 1988; pp. 650)

Dudley and Musgrave (1988) identify two sources of interaction: internal spills and non-linear storage losses. They note that under two restrictive conditions: identical storage decisions (such that internal spills are zero) and linear storage losses, the problem can be condensed to that of a representative agent. They then present a simulation model in which users are assigned policy functions derived from the planners solution.

Alaouze (1991) consider capacity sharing using a simplified analytical framework, in which there is no spot market, no internal spills and linear losses. He demonstrates that capacity sharing achieves equal or better welfare, than an optimal storage but arbitrary (proportional) use allocation scenario. The idea being that tailored storage polices can help minimise water trade requirements.

Recently, Truong and Drynan (2013) presented analytical results for capacity sharing under an assumption of perfect spot markets and no evaporation losses. Under these assumptions capacity sharing achieves a socially optimal outcome in which all users adopt identical storage policies and internal spills never occur (Truong and Drynan 2013).

Brennan (2008; 2010) evaluated government storage policy, using a model of the Goulburn region in Victoria. Brennan (2008) emphasised the role of forfeited (unused) water allocations. Brennan (2008) showed that, given myopic storage policy and an absence of storage rights (our scenario NS), the introduction of trading can decrease welfare, by reducing forfeited allocations and therefore storage reserves.

While Brennan (2008; 2010) makes the case for storage rights she is largely ambivalent about their form. In early work Brennan and Scoccimarro (1999) raised concerns about internal spills under capacity sharing

> ...while the aim of the capacity-sharing institution is to make water users independent of each other the physical reality is that they are interdependent. As an example of such interdependency, the conservative operator will have a large frequency of '[internal] spills' which would increase the volume of water flowing into the capacity shares of the less conservative users in the dam. The management of this water ... has not been dealt with adequately in the literature on capacity sharing (Brennan and Scoccimarro 1999; pp. 84).

Storage rights have also been considered in detail by the Australian Bureau of Agricultural and Resource Economics and Sciences (ABARES). Hughes and Goesch

(2009*b*) outline some of the limitations of standard water rights (i.e., information problems, transfer costs and policy uncertainty) and simple storage rights (i.e., externalities) relative to capacity sharing. Hughes (2010) consider generalising capacity sharing to more complex river systems with multiple storages, inflow sources and demand nodes.

5 Storage rights in the MDB and western US

Currently, all major MDB regulated rivers have some form of user level storage right. Approaches to storage rights can be broadly classified into four systems: 'carryover rights' (Southern NSW), 'spillable water accounts' (Northern VIC), 'continuous accounting' (Northern NSW) and 'capacity sharing' (Southern QLD). Hughes et al. (2013) consider these alternatives in some detail.

A key issue is the annual water accounting framework in the southern MDB, which makes it difficult to internalise storage capacity constraints. In response, the states impose arbitrary rules, such as limits on annual carryover volumes. However, these rules are clumsy and often fail to bind, resulting in close to open access outcomes (Hughes et al. 2013).

A key example, is the Victorian Murray region during the spill events of 2010-11 and 2011-12. During these events storages rights allowed open access despite binding capacity constraints. This resulted in substantial externalities, as users rushed to exploit the situation by transferring water into Victorian storage accounts (Hughes et al. 2013).

Continuous accounting systems involve more frequent water accounting, allowing constraints to be internalised with account limits. While closer to a capacity sharing approach, there remain important differences: account limits don't explicitly match storage capacity, storage losses are socialised, reconciliations occur sporadically and water accounting is monthly (rather than daily as in St George).

In Southern QLD, user level capacity sharing has been adopted in line with the proposals of Dudley and Musgrave (1988). Hughes and Goesch (2009*a*) document the capacity sharing schemes at St George and MacIntyre Brook, observing much enthusiasm for the approach, both among water users and water managers. Water accounting data showed significant heterogeneity in user storage policies and significant, albeit infrequent, internal spills (Hughes and Goesch 2009*a*).

Many of these river systems involve multiple storages. Typically, storage rights are defined over aggregate storage capacity and the problem of distributing water across storages is handled centrally. In most cases, this approximation is likely to be adequate — particularly for storages in series (Hughes 2010). The state level arrangements on the Murray are an exception, where NSW and VIC hold distinct shares to Hume and Dartmouth dams, see MDBA (2011).

Some caution must be taken when linking our policy scenarios with real world systems. Capacity sharing is a broader concept than our CS scenario. Capacity sharing involves reforms taken for granted here, including proportional inflow rights (Hughes et al. 2013). In NSW and VIC, inflows are still allocated centrally, which can lead to policy uncertainty (Hughes et al. 2013). Finally, the Northern Victorian approach is not precisely represented by any single scenario as it combines aspects of the SWA, CS and OA scenarios.

User level storage rights are rare in the Western US. Two exceptions are the Texas Lower-Rio Grande and the South Platte Basin in Colorado, which both have water rights and spot markets more reminiscent of the Southern MDB . Storage rights in the Texas Lower Rio Grande are remarkably similar to the continuous accounting systems of Northern NSW. The Southe Platte basin has a system similar to southern NSW carryover rights.

In central California, ground water banking is typically the first option for storing unused water. However, water contractors (irrigation districts) do hold rights over some storages, including San Louis reservoir. Similar to Northern VIC, San Louis involves 'spill forfeit rules'.

On the Colorado river a form of storage right known as Individually Created Surplus (ICS) emerged in 2007 (Hughes 2013). While these rights are subject to many limitations they effectively allow irrigation districts to store unused water allocations in Hoover dam. Since the introduction of ICS over 1200 GL of water — around 10 per cent of the current balance of Hoover dam — has been accumulated under these rights (Hughes 2013).

6 A parametric model

6.1 Functional form

Inflows I_{t+1} are assumed drawn from a gamma distribution with positive first order autocorrelation (a standard assumption for annual streamflow, see McMahon et al. 2007):

$$I_{t+1} = \rho_I I_t + \epsilon_{t+1}$$

$$\epsilon_{t+1} \sim \Gamma(k_I, \theta_I)$$

$$0 < \rho_I < 1$$

We adopt a standard storage loss function reflective of generic storage geometry (see Lund 2006)

$$L_t = \delta_0.\alpha(S_t)^{2/3}$$

We adopt a linear delivery loss function with fixed and variable losses

$$Q_t \leq \max\{W_t(1-\delta_{1b}) - \delta_{1a}, 0\}$$

The consumptive users $\mathcal{U} = \{1, 2, ..., n\}$ are are grouped into two classes, low reliability users $\mathcal{U}_{low} = \{i \in \mathcal{U} | i \leq n_{low}\}$ (i.e., broadacre farms) and high reliability users $\mathcal{U}_{high} = \{i \in \mathcal{U} | i > n_{low}\}$ (i.e. horticulture farms).

The users have quadratic profit functions $\pi_h(\tilde{q}_{it}, \tilde{I}_t, e_{it})$ where $h \in (low, high)$ indicates membership to $\mathcal{U}_{low}$,$\mathcal{U}_{high}$, $\tilde{I}_t = I_t/E[I_t]$, $\tilde{q}_{it} = q_{it}/\mathscr{A}_h$, $\mathscr{A}_h$ is the farm land area for users in class h and e_{it} is a user specific productivity shock which follows an AR(1) process:

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

$$e_{it} = 1 + \rho_e(e_{i,t-1} - 1) + \eta_{it}$$

$$\eta_{it} \sim \mathcal{N}(0, \sigma_\eta^2)$$

$$0 < \rho_e < 1$$

Here the quadratic expression reflects profit per unit land, conditional on water use per unit land and climate conditions.

6.2 Parameter distributions

In order to maintain generality, parameter distributions are specified rather than point estimates. We provide a brief outline of the parameterisation below. The complete treatment is contained in my draft thesis and can be provided on request.

Supply side parameters are based on a statistical analysis of 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). Where possible, parameter distributions are assumed uniform over the 15th to the 85th percentiles of our data set.

Demand side parameters are based on an econometric analysis of irrigation farms in ABARES survey of irrigation farms in the MDB (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.

Storage capacity K is the numéraire in parameterisation and is fixed at 1000 GL. The central case is defined by the mean values of the parameter distributions. Key parameter assumptions are summarised below:

Table 1: Selected parameter ranges

Figure 2: Social welfare, when $I_t = E[I_t]$ and $e_{it} = E[e_{it}] = 1$.

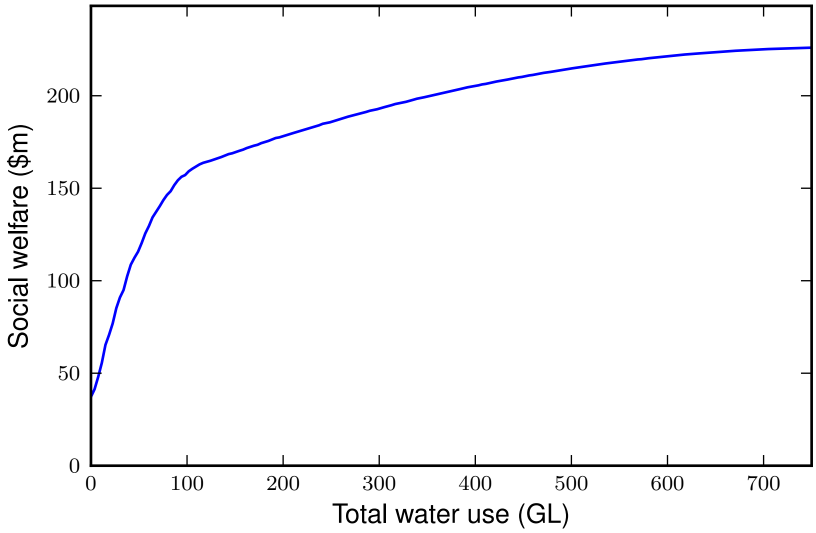

The demand side of the model is summarised by Figure 2, which shows the planners' payoff function implied by the central case parameters — that is $\sum_{\mathcal{U}} u_{it}$ against Q_t, assuming perfect zero transaction cost markets. The steep portion is high reliability demand, which in the central case accounts for around 20 per cent of water use in a full supply year.

The inflow shares λ_i are determined by a single parameter Λ_{high}: the proportion of inflow / storage capacity assigned to high reliability users.

$$\lambda_i = \begin{cases} \Lambda_{high}/n_{high} & \text{if } i \in \mathcal{U}_{high} \\ (1-\Lambda_{high})/n_{low} & \text{if } i \in \mathcal{U}_{low} \end{cases}$$

Inflow shares are set proportional to target demand volumes $\bar{Q}_{high}$ which are defined as $\sum_{\mathcal{U}_{high}} q_{it}$ for an 'average' water year: in which $I_t = E[I_t]$ and $W_t = I_t$. A uniform distribution for Λ_{high} over the range 1 to 2 $\times \bar{Q}_{high}/\bar{Q}$ is selected - since high reliability users tend to demand larger than proportional shares (an issue considered further in my thesis).

7 Solving the model

Given non-market interactions between users, the decentralised model is a stochastic game (Shapley 1953). In stochastic games, each player faces a Markov Decision Process (MDP) where the payoffs and or state transition are dependent on the actions of other players.

Stochastic games present a number of conceptual and practical challenges. The first is defining a solution concept. The second is establishing a method for efficiently computing solutions. We outline our approach below, which is based around the economics of learning in games (Fudenberg and Levine 1998) and the computer science methods of reinforcement learning (Sutton and Barto 1998).

7.1 Equilibrium concepts

In our model, spot market equilibrium is defined by a market clearing price P_t^*, as a function of I_t, $\mathbf{e}_t$ and $\mathbf{w}_t$, which satisfies

$$q_{it} = \tilde{d}_h^{-1}(P_t^*, \tau, a_{it}, I_t, e_{it}) \ \ \forall \ \ i$$

With q_{it}^* determined by the spot market equilibrium, a solution to the users' problem is a policy function for w_{it}

$$w_{it}^* = f_h(\mathbf{s}_t, \mathbf{e}_t, I_t)$$

A Markov Perfect Equilibrium (MPE) (Maskin and Tirole 1988) is then defined by a set of policy functions $f_h(.)$ which simultaneously solve all users' problems. An immediate concern is that the state space for the users' problems scale in the number of users n. With large n this approach is neither feasible nor realistic.

A common response, is to replace opponent state variables with relevant aggregate statistics. Weintraub et al. (2008) describe an equilibrium in these restricted policies as an Oblivious Equilibrium (OE).

Here we assume users have knowledge of the storage volume S_t and inflow I_t, as well as their own account balance s_{it} and productivity shock e_{it} but are 'oblivious' to $\mathbf{s}_t^{-i}$ and $\mathbf{e}_t^{-i}$, restricting our attention to policy functions of the form

$$w_{it}^* = f_h(s_{it}, S_t, e_{it}, I_t)$$

Weintraub et al. (2008) establish existence of OE conditional MPE. Unfortunately, there are no general MPE existence (or uniqueness) results (for a review of the recent literature see Duggan 2012). Weintraub et al. (2008) argue that OE is less susceptible to multiple equilibria in practice.

7.2 Learning in games

The theory of learning in games describes the behaviour of less than fully rational agents in repeated games. In particular, how players adapt their policies in response to observed past play. There is much literature on learning in repeated games, considering how closely different learning models reflect human behaviour and if and when learning models converge on equilibria (see Fudenberg and Levine 1998).

The economic literature on learning in stochastic games is surprisingly scarce. Here more significant contributions have come from the field of computer science. Many recent studies combine computational techniques — such as reinforcement learning — with equilibrium concepts from game theory (for a review see Busoniu et al. 2008). Fudenberg and Levine (2007) provide an economic perspective on this literature.

7.3 A reinforcement learning approach

Reinforcement learning (also known as approximate and neuro 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 — based on the method of 'Fitted Q iteration' (Ernst et al. 2005) — is summarised briefly below. To begin we derive guesses for user policy $\hat{f}_h(\mathscr{X}_{it})$ and value functions $\hat{v}_h(\mathscr{X}_{it})$ from the planners solution, where $\mathscr{X}_{it}$ is the state vector (s_t, S_t, e_t, I_t). Then the users' problems are solved — holding opponent polices fixed at initial guesses — by fitted Q-V iteration. Then we proceed to a full learning algorithm, where the population of polices varies between each iteration, this stage is outlined below

7.3.1 Multiple agent fitted Q-V iteration

1. Simulate the decentralised storage problem for T periods, using current user policy functions $\hat{f}_i(\mathscr{X}_{it})$, with exploration polices (i.e., partially randomised

policies) assigned to a small subset of users $\mathcal{U}^e \subset \mathcal{U}$. Obtain the action w_{it}, payoff u_{it} and state $\mathscr{X}_{it} = (s_{it}, S_t, e_{it}, I_t)$ samples

$$\{w_{it}, u_{it}, \mathscr{X}_{it}, \mathscr{X}_{it+1} | t = 1, ..., T, i \in \mathcal{U}^e\}$$

2. Generate the Q 'action-value' function samples q_{it} for $t = 1, ..., T$, $i \in \mathcal{U}^e$

$$q_{it} \leftarrow u_{it} + \beta \hat{v}_h(\mathscr{X}_{it+1})$$

3. Fit continuous Q functions $\hat{Q}_h(w_{it}, \mathscr{X}_{it})$ for each user class. This step involves regression problems, with dependent variables q_{ht} and explanatory variables w_{ht}, $\mathscr{X}_{ht}$ (with samples grouped by user class h).

4. Optimise the Q functions for a sub set of the state samples and fit updated policy $\hat{f}_h^1(.)$ and value $\hat{v}_h^1(.)$ functions

$$\hat{f}_h^1(\mathscr{X}_{it}) = \arg\max_{w_{it} \leq s_{it}} \hat{Q}_h(w_{it}, \mathscr{X}_{it})$$

$$\hat{v}_h^1(\mathscr{X}_{it}) = \max_{w_{it} \leq s_{it}} \hat{Q}_h(w_{it}, \mathscr{X}_{it})$$

5. Assign updated policy functions $\hat{f}_h^1$ to a random sample of users $\mathcal{U}^1 \subset \mathcal{U}$, then return to step 1 and repeat for a fixed number of iterations.

The above approach can be interpreted as a learning method. Within the computer science literature the approach represents 'rational' agent learning (Bowling and Veloso 2001): learning that converges on best response policies given stationary opponent policies. Within the economic literature, the approach might be described as an 'optimisation-based' learning method (Crawford 2013).

While the approach is to be interpreted as a learning method, it represents only a small departure from algorithms used to computed rational expectations equilibria. For example, the approach is similar to the value iteration method used to solve stochastic games (Shapley 1953). It is also related to the Krusell and Smith (1998) style algorithms used to solve macro heterogeneous agent models. In many ways, the approach represents a middle ground between the rational expectations methods of modern macroeconomics and the simulation and genetic algorithm methods of agent based modeling (Tesfatsion and Judd 2006).

7.4 Computation

Successful implementation of reinforcement learning rests on algorithm design choices particularly: sample sizes, exploration policies and function approximation schemes. A complete discussion of the computational approach — contained in my draft thesis — can be provided on request.

For function approximation we use a version of tile coding (Sutton and Barto 1998). The model is coded in python. Time sensitive components are translated to compiled c via cython. The implementation makes extensive use of parallel computing. The whole process completes in under 5 minutes on a standard 4 core desktop.

8 Results

8.1 Central case

We begin with results for the central case parameters.

Each iteration of the algorithm involves a simulation of length T. To begin, we show how the sample means from each simulation stage evolve as the learning algorithm progresses. For example, Figure 3 shows the mean storage level $\frac{1}{T}\sum_{t=1}^{T} S_t$ at each iteration — beginning with the planners solution as iteration 0. Our final results are sample statistics from the last iteration of the algorithm, these are presented in Tables 3 to 7.

Figure 3: Mean storage by iteration

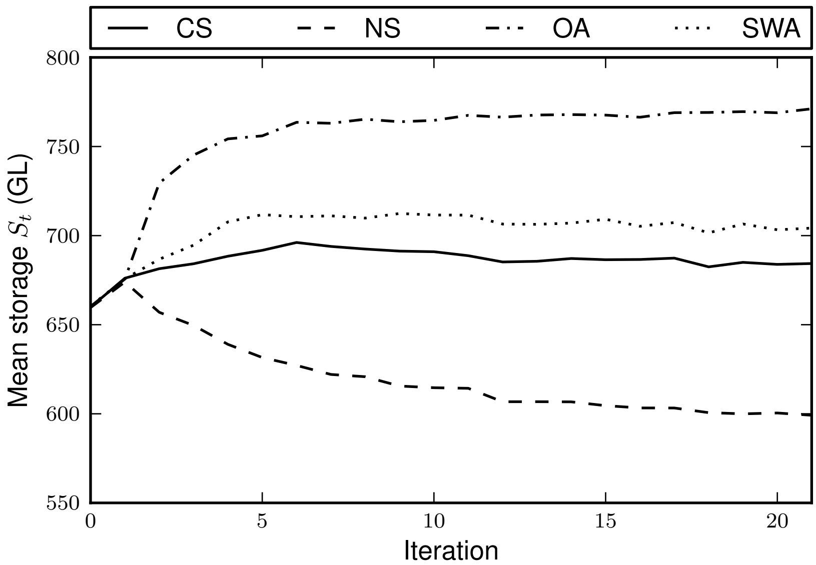

Immediately some expected results are apparent. Open access results in significant over storage and no storage access in significant under storage. Spill forfeit rules generate higher storage levels than capacity sharing — which is closest to the planner's solution.

Importantly the differences between scenarios are stable over the course of the algorithm. While the algorithm does not converge to a precise equilibrium, the user value and policy functions show a tendency to converge rather than diverge or cycle spectacularly (Figures 5 and 6).

Figure 4: Mean storage reserve by iteration

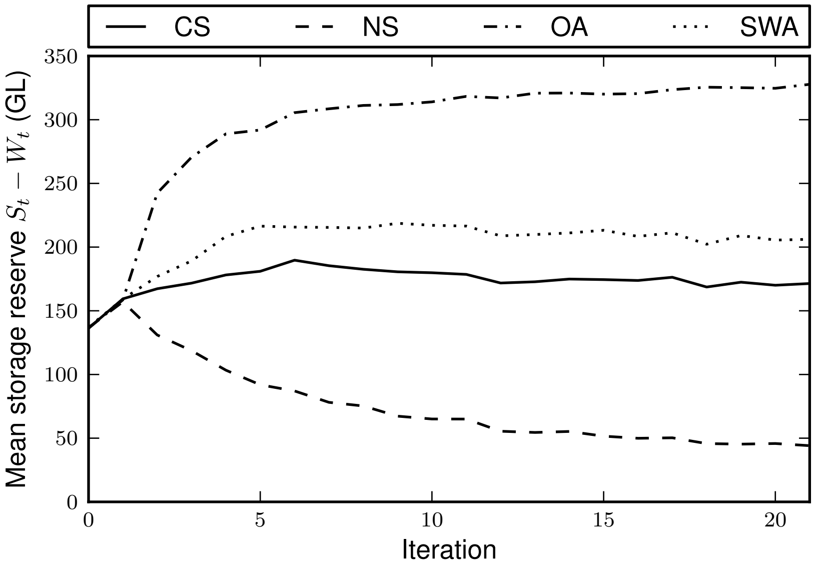

Figure 5: Value function error (Mean absolute percentage deviation)

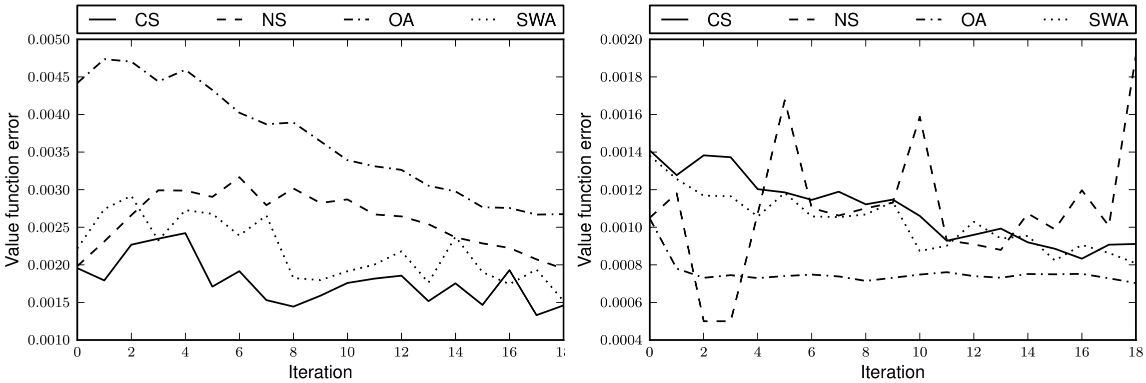

(a) Low reliability

(b) High reliability

Figure 6: Policy function error (Mean absolute percentage deviation)

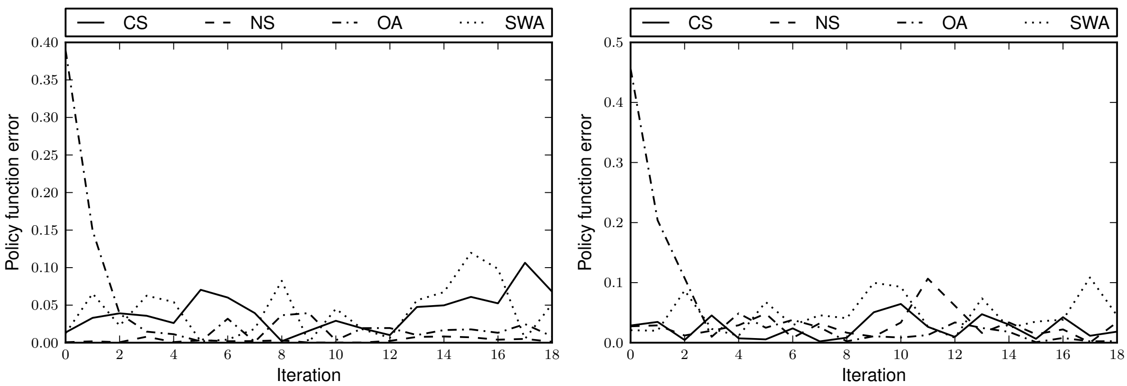

(a) Low reliability

(b) High reliability

In the central case, capacity sharing results in the highest mean social welfare, followed closely by spill forfeit rules (Figure 7, Table 2). Open access storage is only a modest improvement over no storage access.

Figure 7: Mean social welfare by iteration

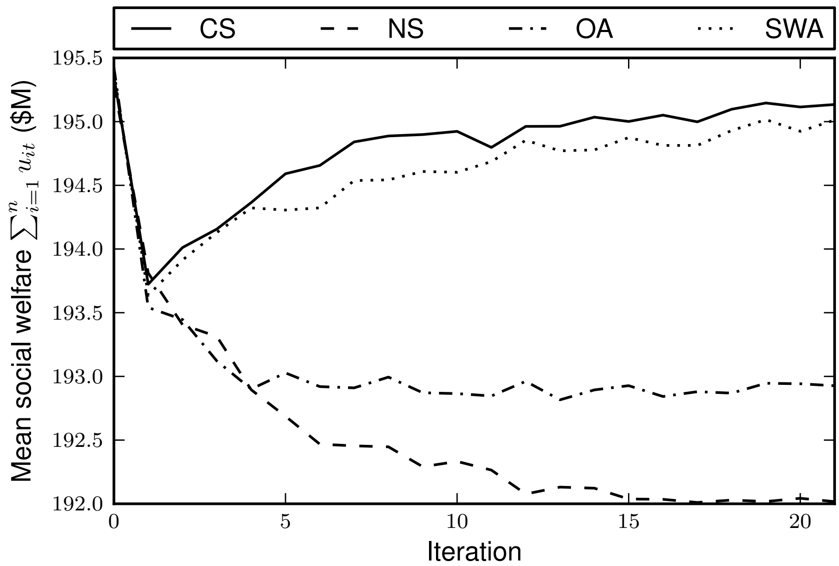

Unsurprisingly, the mean welfare differences are relatively small — at least in this central case. As is standard in the storage literature, changes in the variance of welfare are larger than changes in the mean (see Table 2). Further, the aggregate welfare effects hide some larger distributional results.

Scenarios that result in over storage (OA, SWA) favour high reliability users at the expense of low reliability users — and visa versa (see Tables 6 and 7). These effects are explained partly by the storage differences, but also by the nature of the externalities x_{it} (discussed below, see Figure 10)

Figure 8 shows mean withdrawals W_t (the sum of user withdrawals) conditional on the storage level. Note that, mean withdrawals are slightly lower in spill years (when $S_t = K$) as demand for water is depressed by high inflows.

Under all scenarios, a significant degree of heterogeneity is observed in user storage policies. Figure 9 shows mean withdrawals as a proportion of account levels for high and low reliability user groups. Even with a relatively moderate transaction cost, we observe a high degree of specialisation: high reliability users adopt more conservative storage policy.

Figure 8: Aggregate withdrawal policy, $E[W_t|S_t]$

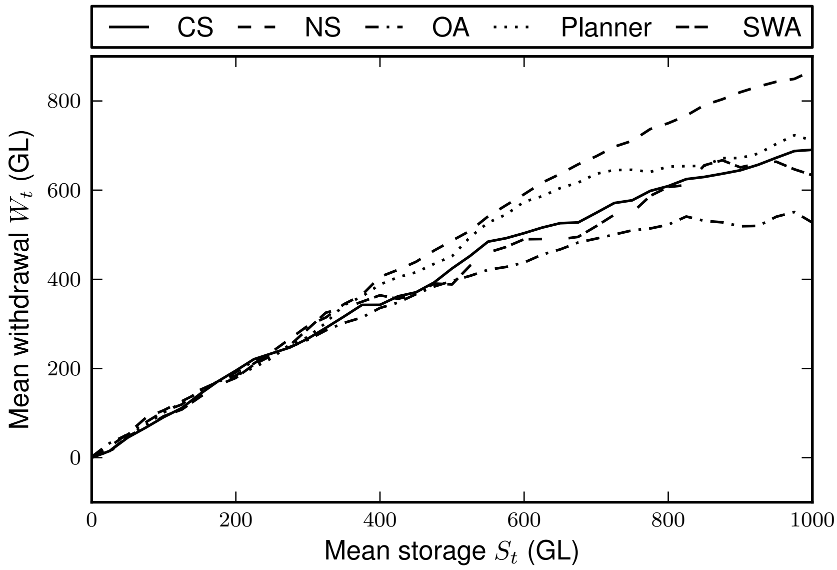

Figure 9: User withdrawals over storage, $\frac{\frac{1}{T}\sum_{t=1}^{T}\sum_{\mathcal{U}^h} w_{it}}{\frac{1}{T}\sum_{t=1}^{T}\sum_{\mathcal{U}^h} s_{it}}$

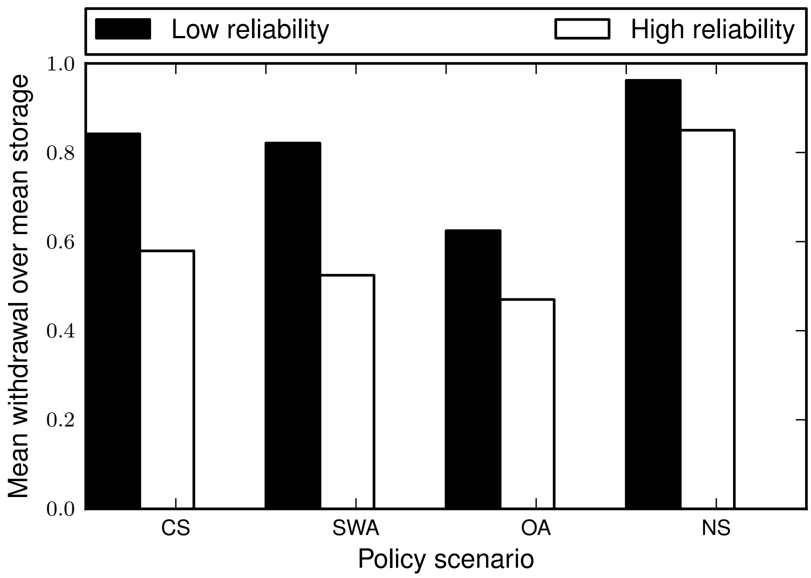

Figure 10 shows mean externalities as a proportion of mean account balances for low and high user groups. Externalities are highest under OA and lowest under CS. Under CS, externalities are positive due to internal spills. Internal spills tend to favour low reliability users, as they flow from users with high balances to those with low balances. In contrast, externalities favour high reliability users under open access.

Figure 10: User externalities over storage, $\frac{\frac{1}{T}\sum_{t=1}^{T}\sum_{\mathcal{U}^h} x_{it}}{\frac{1}{T}\sum_{t=1}^{T}\sum_{\mathcal{U}^h} s_{it}}$

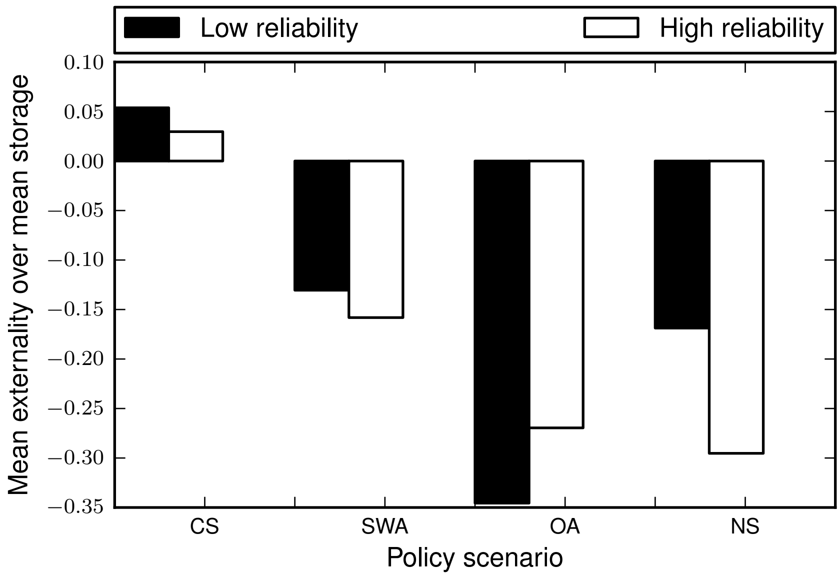

Tables 3 to 7 also contain results for the socialised evaporation loss scenarios (CS-SL, SWA-SL). Socialising losses leads to small increases in mean storage and decreases in mean welfare. The effect of socialising losses is more pronounced under the SWA scenario. On a distributional level, socialised losses tend to favour high reliability users.

Finally, the large differences in storage levels between scenarios, lead to some significant changes in storage spills (see Table 5) which may have welfare implications for in-stream users — or for downstream consumptive users.

Table 2: Social welfare $\sum_{i=1}^{n} u_{it}$ ($M)

	Mean	SD	2.5th	25th	75th	97.5th
CS	195.1	22.8	136.1	188.0	209.4	217.4
CS-SL	195.1	21.4	142.5	188.6	208.4	216.4
SWA	195.0	22.0	136.7	189.6	208.5	216.4
SWA-SL	194.7	20.0	143.5	190.8	206.4	213.4
OA	192.9	17.6	148.5	189.2	202.7	209.7
NS	192.0	27.8	115.0	179.9	210.9	218.8
Planner	195.4	26.1	124.5	188.1	211.5	219.3

Table 3: Storage S_t (GL)

	Mean	SD	2.5th	25th	75th	97.5th
CS	684.3	287.8	133.1	445.8	1,000.0	1,000.0
CS-SL	701.7	281.8	145.7	473.4	1,000.0	1,000.0
SWA	704.2	287.3	134.4	467.9	1,000.0	1,000.0
SWA-SL	729.3	278.6	148.3	511.2	1,000.0	1,000.0
OA	771.2	269.3	161.8	580.9	1,000.0	1,000.0
NS	599.2	302.0	102.3	339.0	923.2	1,000.0
Planner	659.8	298.6	114.3	404.3	1,000.0	1,000.0

Table 4: Withdrawal W_t (GL)

	Mean	SD	2.5th	25th	75th	97.5th
CS	512.9	174.0	126.8	386.5	669.2	693.1
CS-SL	505.2	162.5	138.4	391.0	649.4	672.5
SWA	498.1	152.6	127.9	402.0	615.7	650.6
SWA-SL	483.2	131.9	140.9	411.7	578.9	599.2
OA	443.4	99.1	154.6	406.9	512.9	544.4
NS	555.1	259.6	97.7	331.6	831.4	886.8
Planner	523.5	190.6	114.3	382.9	676.4	721.5

Table 5: Spills Z_t (GL)

	Mean	SD	2.5th	25th	75th	97.5th
CS	123.9	272.2	0.0	0.0	67.2	1,040.4
CS-SL	131.4	280.5	0.0	0.0	93.0	1,066.0
SWA	138.7	288.8	0.0	0.0	115.0	1,077.3
SWA-SL	150.9	301.0	0.0	0.0	152.4	1,119.5
OA	189.9	336.2	0.0	0.0	256.0	1,228.5
NS	87.9	225.7	0.0	0.0	0.0	960.4
Planner	115.3	262.0	0.0	0.0	36.2	993.2

Table 6: Total low reliability payoff $\sum_{i \in \mathcal{U}^{low}} u_{it}$ ($M)

	Mean	SD	2.5th	25th	75th	97.5th
CS	85.5	13.1	57.3	76.0	96.1	100.7
CS-SL	85.0	12.8	56.7	76.1	95.3	99.9
SWA	84.9	12.6	56.6	76.0	94.7	99.2
SWA-SL	84.0	11.5	57.0	76.8	92.6	96.7
OA	82.2	8.6	59.0	79.8	87.7	92.1
NS	86.3	13.1	61.5	75.6	98.0	103.0
Planner	83.8	19.2	42.0	72.9	97.7	102.3

Table 7: Total high reliability payoff $\sum_{i \in \mathcal{U}^{high}} u_{it}$ ($M)

	Mean	SD	2.5th	25th	75th	97.5th
CS	109.6	15.1	62.0	107.7	116.4	120.9
CS-SL	110.2	13.5	74.9	108.0	116.3	120.8
SWA	110.2	15.4	61.4	108.3	117.0	121.5
SWA-SL	110.7	13.8	74.4	109.1	116.7	121.1
OA	110.8	12.7	85.8	109.1	116.3	120.6
NS	105.8	20.1	32.6	104.1	116.0	121.1
Planner	111.5	12.4	82.2	110.7	116.6	120.5

8.2 General case

Here 110 parameter sets were randomly drawn. For each set of parameters the model was solved for the CS, SWA, OA and NS scenarios. For each parameter set and each scenario, we calculate the following sample means (from the final iteration of the learning algorithm):

- Mean social welfare: $\frac{1}{T}\sum_{t=1}^{T}\sum_{i=1}^{n} u_{it}$
- Mean low and high reliability welfare: $\frac{1}{T}\sum_{t=1}^{T}\sum_{i\in\mathcal{U}^{low}} u_{it}$, $\frac{1}{T}\sum_{t=1}^{T}\sum_{i\in\mathcal{U}^{high}} u_{it}$
- Mean storage: $\frac{1}{T}\sum_{t=1}^{T} S_{it}$
- Mean spills: $\frac{1}{T}\sum_{t=1}^{T} Z_{it}$

For each statistic we also define an index relative to the CS scenario: Y_j^m/Y_j^{CS}, where Y_j^m is the value of statistic Y for scenario $m \in \{CS, SWA, OA, NS\}$ under parameter set j. Summary statistics over these sample means and their indexes, are presented at the end of this section in Tables 8 to 19

8.2.1 Social welfare

Social welfare results are summarised in Tables 8 and 9 and Figure 11.

Figure 11: Social welfare index

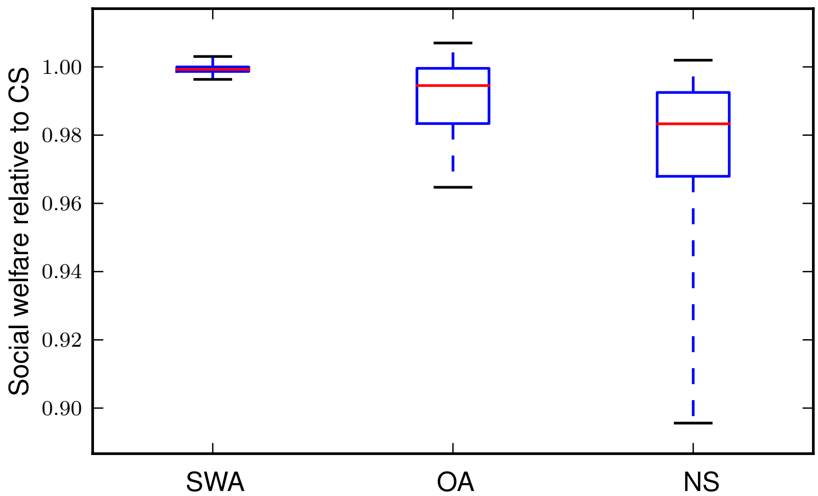

On average, CS achieves the highest mean welfare. On social welfare grounds, CS is the preferred scenario in 65 of the 110 cases, OA in 20, SWA in 16 and NS in five. In those cases where CS is not preferred, the welfare differences are small. The welfare differences between SWA and CS are almost always trivial. By far the largest welfare losses occur in the NS scenario.

We can gain further insight by examining correlations between our indexes and the parameters. Here the mean welfare indexes were regressed against the parameters, using a non-parametric Random Forests method (Geurts et al. 2006).

The two most important parameters are: the ratio of mean inflow to storage capacity and inflow variation (Table 11). In low inflow (and high variance) cases NS performs relatively poorly (see Figure 12). The inverse result is observed for the OA scenario — and to a much lesser extent the SWA scenario.

The intuition here is that in high inflow cases the cost of storage (i.e., the risk of spill) is high relative to the benefits. As such OA — which socialises these costs — performs poorly, while NS — which socialises both benefits and costs — performs relatively well.

Figure 12: Welfare index regression results

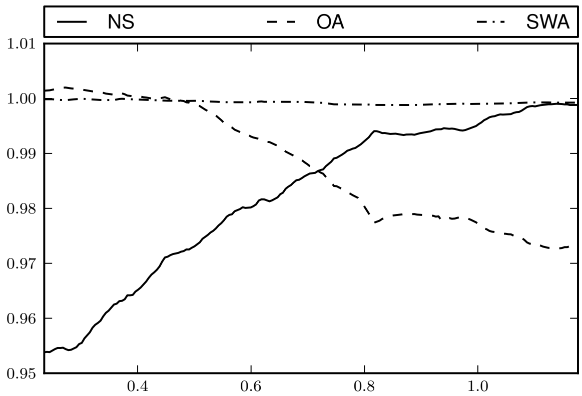

(a) Mean inflow to capacity, $E[I_t]/K$

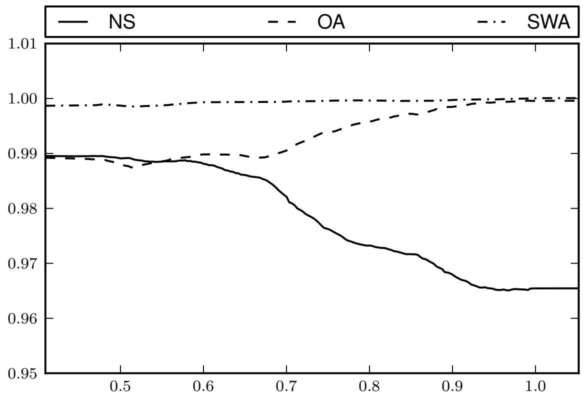

(b) Inflow variability, cv_I

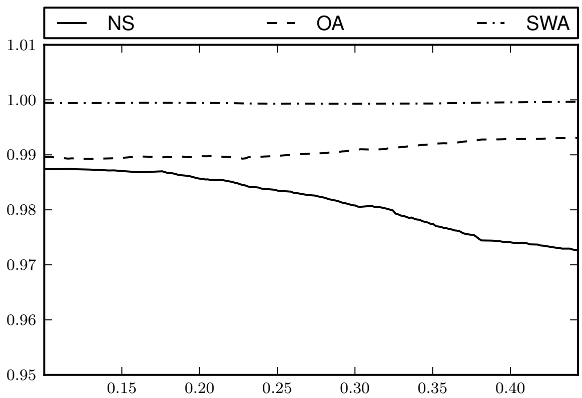

(c) Mix of high demand, $\bar{Q}^{high}/\bar{Q}$

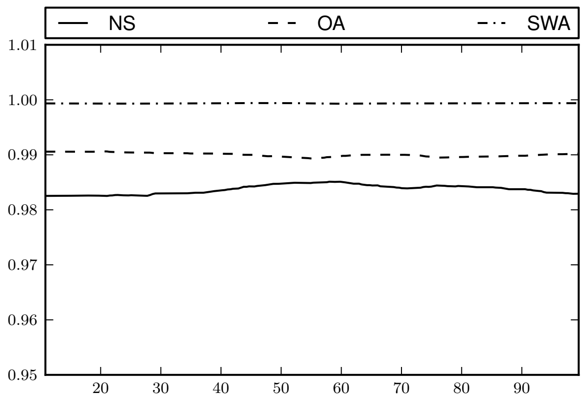

(d) Transaction cost, τ

Figure 13 plots the preferred scenario against the two main inflow parameters. Here the prefered scenario was also regressed (as a qualitative dependent variable) against the parameters (using a Random Forest classifier). Figure 13 shows the predicted space in which each scenario is preferred (given other parameters fixed at sample means). Specifically, Figure 13 shows that CS is the preferred scenario for the majority of the parameter space (the shaded red area), but the OA is preferred in some low inflow / high variance cases (the shaded blue area).

The transaction cost τ, and the mix of high / low users also have significant effects on the preferred scenario (Table 10): higher transaction costs favour SWA over CS, because internal spills become more significant.

Figure 13: Preferred scenario by inflow parameters, with classifier results

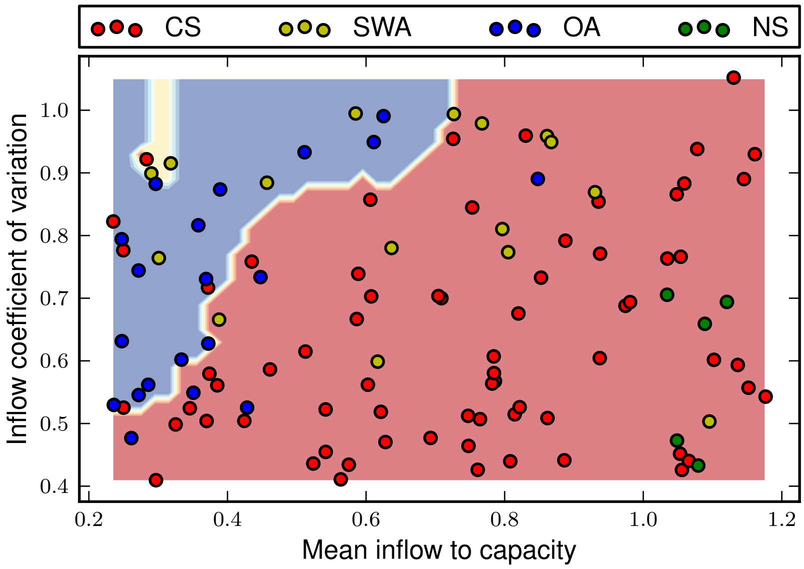

8.2.2 Welfare distribution

Distributional differences are generally larger than social welfare differences, see Tables 12, 13, 14 and 15 and Figures 14 and 15 below. In general, OA favours high reliability users, while NS favours low reliability users. While SWA and CS are barely separable on social welfare grounds, there are some noticeable distributional differences, with SWA favouring high reliability users.

Figure 14: Low reliability payoff index

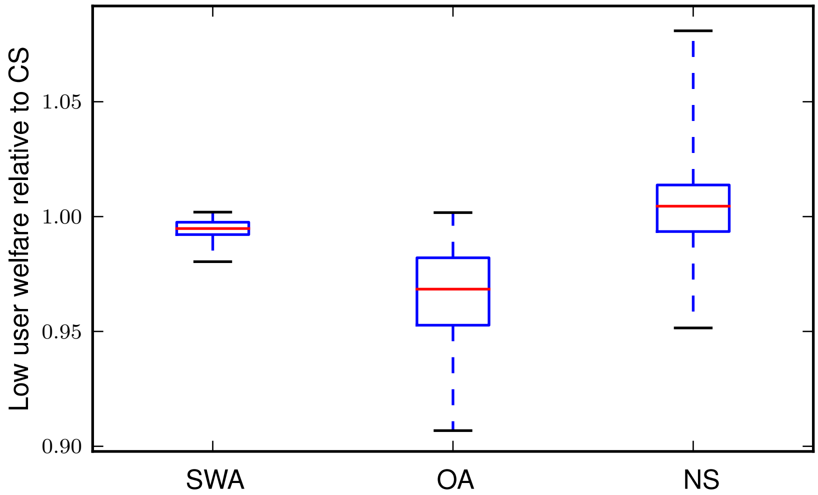

Figure 15: High reliability payoff index

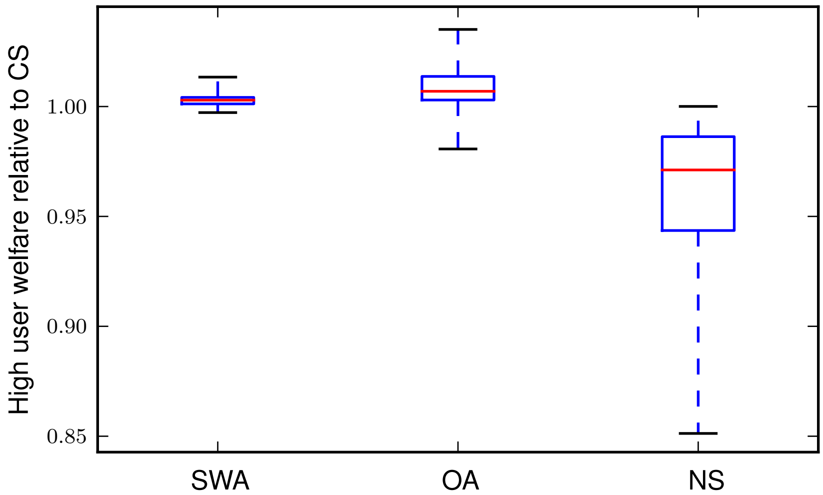

8.2.3 Storage

The scenarios all induce significant changes in mean storage levels (Tables 16 and 17 and Figure 16). In almost all cases, OA induces significantly higher storage reserves than CS, NS significantly lower and SWA slightly higher.

Figure 16: Storage index

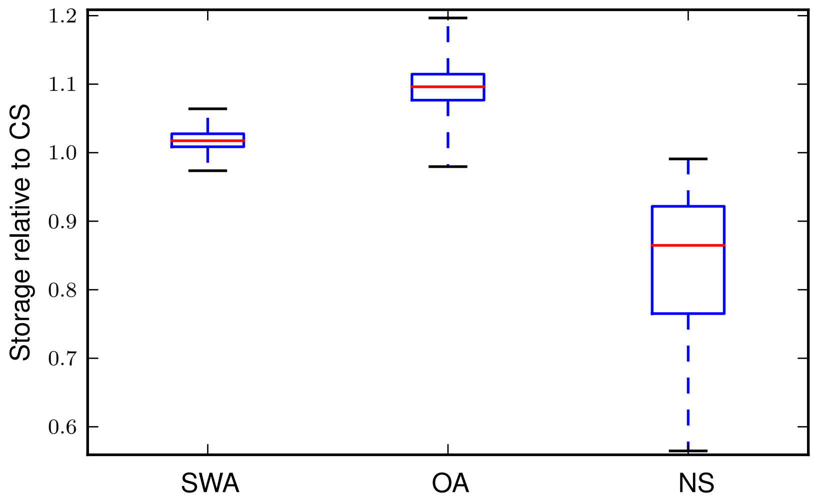

8.2.4 Spills

Changes in mean storage levels lead to amplified changes in storage spills (Tables 18 and 19 and Figure 17). Higher mean spills reflect both an increase in the frequency and magnitude of spill events.

Figure 17: Spills index

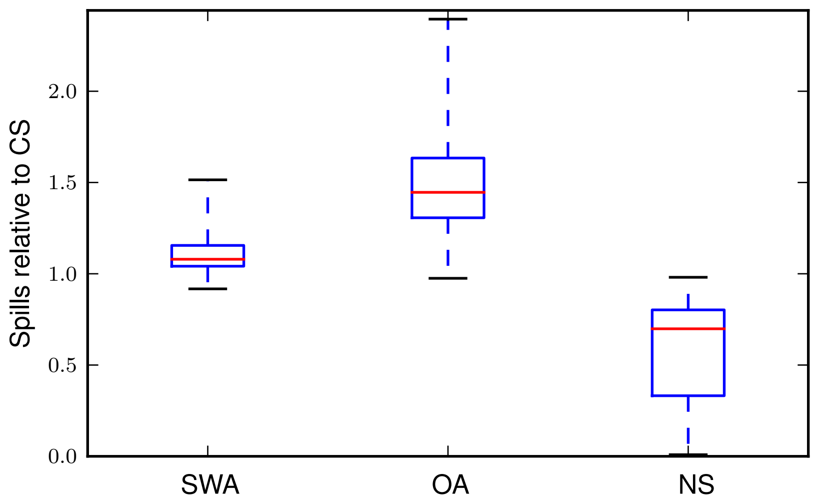

Table 8: Mean social welfare $\frac{1}{T}\sum_{t=1}^{T}\sum_{i=1}^{n} u_{it}$ ($M)

	Mean	Min	Q1	Q3	Max
CS	181.90	45.86	114.67	235.16	370.16
SWA	181.76	45.85	114.40	234.92	369.91
OA	179.73	45.78	114.59	230.27	362.79
NS	178.93	43.05	110.08	232.94	366.05
Planner	182.68	46.45	114.08	235.23	372.18

Table 9: Social welfare index

	Mean	Min	Q1	Q3	Max
CS	1.000	1.000	1.000	1.000	1.000
SWA	0.999	0.996	0.999	1.000	1.003
OA	0.991	0.965	0.983	1.000	1.007
NS	0.978	0.896	0.968	0.993	1.002

Table 10: Preferred scenario classifier: parameter importance and sample means

	Importance	CS	SWA	OA	NS
$E[I]/K$	38.44	0.74	0.65	0.39	1.07
cv_I	28.61	0.64	0.83	0.72	0.59
τ	14.53	51.63	66.85	50.77	38.62
$\bar{Q}_{high}/\bar{Q}$	13.83	0.23	0.22	0.20	0.23

Table 11: Welfare index regression: parameter importance

	Importance
$E[I]/K$	74.00
cv_I	22.28
$\bar{Q}_{high}/\bar{Q}$	3.16
τ	0.47
$\bar{P}$	0.10

Table 12: Mean low reliability payoff $\frac{1}{T}\sum_{t=1}^{T}\sum_{i\in\mathcal{U}^{low}} u_{it}$ ($M)

	Mean	Min	Q1	Q3	Max
CS	76.30	13.88	44.58	106.09	162.84
SWA	75.85	13.88	44.32	105.78	162.53
OA	73.47	13.80	43.55	102.71	156.16
NS	76.80	13.73	45.51	107.84	163.58

Table 13: Low reliability payoff index

	Mean	Min	Q1	Q3	Max
CS	1.000	1.000	1.000	1.000	1.000
SWA	0.994	0.980	0.992	0.998	1.002
OA	0.967	0.907	0.953	0.982	1.002
NS	1.005	0.952	0.993	1.014	1.081

Table 14: Mean high reliability payoff $\frac{1}{T}\sum_{t=1}^{T}\sum_{i\in\mathcal{U}^{high}} u_{it}$ ($M)

	Mean	Min	Q1	Q3	Max
CS	105.60	15.52	60.25	143.37	254.08
SWA	105.91	15.51	60.54	144.10	255.63
OA	106.26	15.66	60.88	145.66	255.06
NS	102.13	14.74	57.12	138.33	245.58

Table 15: High reliability payoff index

	Mean	Min	Q1	Q3	Max
CS	1.000	1.000	1.000	1.000	1.000
SWA	1.003	0.997	1.001	1.004	1.013
OA	1.008	0.981	1.003	1.014	1.035
NS	0.962	0.851	0.944	0.986	1.000

Table 16: Mean storage $\frac{1}{T}\sum_{t=1}^{T} S_t$ (GL)

	Mean	Min	Q1	Q3	Max
CS	635.78	276.09	532.95	744.67	857.13
SWA	648.07	274.30	547.83	760.72	874.25
OA	697.28	318.34	594.96	823.08	908.92
NS	546.44	229.80	391.17	674.30	830.27
Planner	631.27	325.34	538.49	723.27	830.71

Table 17: Storage index

	Mean	Min	Q1	Q3	Max
CS	1.000	1.000	1.000	1.000	1.000
SWA	1.018	0.974	1.009	1.028	1.064
OA	1.096	0.980	1.077	1.115	1.197
NS	0.843	0.565	0.765	0.922	0.991

Table 18: Mean spills $\frac{1}{T}\sum_{t=1}^{T} Z_t$ (GL)

	Mean	Min	Q1	Q3	Max
CS	114.22	0.20	27.49	195.81	302.45
SWA	123.91	0.19	30.21	212.45	315.96
OA	166.53	0.47	31.47	269.29	423.37
NS	87.40	0.01	9.56	154.58	243.69
Planner	107.84	0.42	28.37	184.11	293.43

Table 19: Spills index

	Mean	Min	Q1	Q3	Max
CS	1.000	1.000	1.000	1.000	1.000
SWA	1.106	0.917	1.041	1.155	1.515
OA	1.498	0.975	1.307	1.634	2.395
NS	0.592	0.009	0.332	0.802	0.981

9 Conclusions

9.1 Research questions

After reviewing the results we can address our initial questions:

Which system of storage rights maximises social welfare?

The simple answer is capacity sharing (i.e., storage capacity rights with loss deductions). The more complex answer is: it depends. CS is the most frequently preferred scenario, but each scenario is preferred in at least some cases. On social welfare grounds the difference between storage capacity rights and spill forfeit rules is trivial. However, in many cases open access and no storage rights generate significant welfare losses by inducing storage behaviour that is far from optimal.

The social welfare effects depend to a large extent on the ratio of inflow to storage capacity. In river systems with low inflow to storage capacity, the welfare costs of open access (no storage rights) are lower (higher).

What are the distributional effects?

Storage right systems can have significant effects on the distribution of welfare between low and high reliability user classes. Open access favours high reliability users and no storage rights favours low reliability users. Spill forfeit rules favour high reliability users in comparison with storage capacity rights.

What are the effects on storage levels?

While social welfare effects between scenarios are sometimes trivial, storage effects are often large. Open access results in substantial over storage, while no storage rights results in substantial under storage. Spill forfeit rules result in non-trivial increases in mean storage levels relative to storage capacity rights.

These changes in storage levels lead to amplified changes in spills, which may have implications for downstream users or for in-stream values, particularly flood mitigation or environmental flows. Since central storage release rules are typically on the aggressive side (Brennan 2008, Hughes and Goesch 2009*b*), transitioning from no storage rights to a system of storage rights (whether it be CS, SWA or OA) is likely to lead to an increase in mean storage levels and spills.

9.2 Policy implications

In many cases the welfare differences between scenarios are trivial and will likely be outweighed by transition costs. In other cases, a transition from no storage access to some form of storage right may offer a significant gain. However, much of this gain can go unrealised, if storage rights approximate an open access outcome.

Two analogies can be drawn between our findings and some well known natural resource results. Firstly, the 'Gisser-Sanchez Effect': that the welfare gains gains from optimal groundwater extraction are often (but not always) trivial (Koundouri 2004). Secondly, the idea of limited-user 'open access' fisheries (Wilen 1979): where governments establish quota systems but set non-binding catch limits — incurring the costs of regulation without the benefits.

The preferred approach to storage rights will depend greatly on the river system. Given our results, it is understandable that capacity sharing has been implemented in Northern MDB (where inflows are high relative to storage) and spill forfeit rules in the south (where spills are less frequent). While existing approaches may be adapted to local conditions, recent developments — particularly predictions of lower and more variable inflows under climate change – provide grounds for some reconsideration.

The central conclusion from this study is that, where well implemented, spill forfeit rules or storage capacity rights, can produce a close to optimal outcome. That is, the externalities they generate — while relevant for the distribution of welfare and for storage levels — have trivial effect on social welfare. This conclusion may change however, in the case where storage spills have welfare effects.

References

Alaouze, C. M. 1991, 'The optimallity of capacity sharing in stochastic dynamic programing problems of shared reservoir operation', *Journal of The American Water Resources Association*, vol. 27, pp. 381–386.

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>

Bowling, M. and Veloso, M. 2001, Rational and convergent learning in stochastic games, *in* 'International joint conference on artificial intelligence', Vol. 17, pp. 1021–1026.

Brennan, D. 2008, 'Missing markets for storage and the potential economic cost of expanding the spatial scope of water trade', *Australian Journal of Agricultural and Resource Economics*, vol. 52, pp. 471–485.

Brennan, D. 2010, 'Economic potential of market-oriented water storage decisions: Evidence from Australia', *Water Resources Research*.

Brennan, D. and Scoccimarro, M. 1999, 'Issues in defining property rights to improve Australian water markets', *Australian Journal of Agricultural and Resource Economics*, vol. 43, pp. 69–89.

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

Busoniu, L., Babuska, R. and De Schutter, B. 2008, 'A comprehensive survey of multiagent reinforcement learning', *Systems, Man, and Cybernetics, Part C: Applications and Reviews, IEEE Transactions on*, vol. 38, no. 2, pp. 156–172.

Crawford, V. P. 2013, 'Boundedly Rational versus Optimization-Based Models of Strategic Thinking and Learning in Games', *The Journal of Economic Literature*, vol. 51, no. 2, pp. 512–27.

Dudley, N. 1992, 'Water allocation by markets, common property and capacity sharing: companions or competitors', *Natural Resources Journal*, vol. 32, pp. 757.

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

Dudley, N. J. 1988*b*, 'Volume sharing of reservoir water', *Water Resources Research*, vol. 24, no. 5, pp. 641–648.

Dudley, N. J. 1999, 'Water resource sharing from a microeconomic perspective', *Cambridge Review of International Affairs*, vol. 12, no. 2, pp. 239–253.

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

Duggan, J. 2012, 'Noisy stochastic games', *Econometrica*, vol. 80, no. 5, pp. 2017–2045.

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.

Fudenberg, D. and Levine, D. K. 2007, 'An economist's perspective on multi-agent learning', *Artificial intelligence*, vol. 171, no. 7, pp. 378–381.

Geurts, P., Ernst, D. and Wehenkel, L. 2006, 'Extremely randomized trees', *Machine learning*, vol. 63, no. 1, pp. 3–42.

Hughes, N. 2010, Defining property rights to water in complex regulated river systems: generalising the capacity sharing concept, *in* '54th Annual Australian Agricultural and Resource Economics Society Conference', Adelaide, Australia.

Hughes, N. 2013, 'The water storage problem'. <http://www.globalwaterforum.org/2013/12/09/the-water-storage-problem/>

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.

Hughes, N. and Goesch, T. 2009*a*, *Capacity sharing in the St George and MacIntyre Brook irrigation schemes in southern Queensland*, Australian Bureau of Agricultural and Resource Economics.

Hughes, N. and Goesch, T. 2009*b*, *Management of irrigation water storages: carryover rights and capacity sharing*, Australian Bureau of Agricultural and Resource Economics.

Koundouri, P. 2004, 'Potential for groundwater management: Gisser-Sanchez effect reconsidered', *Water Resources Research*.

Krusell, P. and Smith, J. A. A. 1998, 'Income and wealth heterogeneity in the macroeconomy', *Journal of Political Economy*, vol. 106, no. 5, pp. 867–896.

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.

Maskin, E. and Tirole, J. 1988, 'A theory of dynamic oligopoly, I: Overview and quantity competition with large fixed costs', *Econometrica: Journal of the Econometric Society*, vol. pp. 549–569.

McMahon, T. A., Vogel, R. M., Peel, M. C. and Pegram, G. G. 2007, 'Global streamflows–Part 1: Characteristics of annual streamflows', *Journal of Hydrology*, vol. 347, no. 3, pp. 243–259.

Murray-Darling Basin Authority 2011, *Water sharing in the River Murray System including resources assessment, system operations and bulk water accounting*.

Murray-Darling Basin Ministerial Council 1995, *An Audit of Water Use in the Murray-Darling Basin - June 1995*.

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

Randall, A. 1983, 'The Problem of Market Failure', *Natural Resources Journal*, vol. 23, pp. 131–148.

Shapley, L. 1953, 'Stochastic games', *Proceedings of the National Academy of Sciences of the United States of America*, vol. 39, no. 10, pp. 1095.

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

Tesfatsion, L. and Judd, K. L. 2006, *Agent-based computational economics*, Vol. 2.

Truong, C. H. and Drynan, R. G. 2013, 'Capacity sharing enhances efficiency in water markets involving storage', *Agricultural Water Management*, vol. 122, pp. 46–52.

Weintraub, G., Benkard, C. and Van Roy, B. 2008, 'Markov Perfect Industry Dynamics With Many Firms', *Econometrica*, vol. 76, no. 6, pp. 1375–1411.

Wilen, J. E. 1979, 'Fisherman behavior and the design of efficient fisheries regulation programs', *Journal of the Fisheries Board of Canada*, vol. 36, no. 7, pp. 855–858.