Water property rights in rivers with large dams

Neal Hughes

A thesis submitted for the degree of Doctor of Philosophy of the Australian National University

October 2015

Declaration

I declare that this thesis is my own original work. This thesis has not been published elsewhere or used for the qualification of any other degree or diploma.

Neal Hughes

October 2015

For Elizabeth. Ta-da!

Acknowledgements

I would like to thank the Sir Roland Wilson Foundation for their Phd scholarship. This thesis would not have been possible without their support. I would also like to thank my employer the Australian Government Department of Agriculture / ABARES for supporting my work and for providing access to data.

This work has benefited greatly from the input of my academic supervisors at the ANU: Quentin Grafton, John Stachurski and Sara Beavis.

Over the last three years I've also benefited from conversations with: Steve Beare (Analytecon), Peter Gooday and others at ABARES, John Williams (ANU), Mac Kirby (CSIRO), David Galeano and Jim Foreman (MDBA), Dermot McKane (BOM), John Freebairn (University of Melbourne), John Quiggin (University of Queensland), Andrew Ross (ANU) and Simon Hone (Latrobe University).

I'd also like to thank Julia Jasonsmith (Murrang Earth Sciences) and others from the AWA ACT branch for their feedback on chapter 5 and Peter Sunehag and Scott Sanner and others from the ANU Research School of Computer Science for their feedback on chapter 8.

I'm grateful to Pat Hulme (Sustainable Soil Management) for providing data on Australian irrigation system losses and Murray Peel and Thomas McMahon (University of Melbourne) for providing statistics on Australian rivers.

This work benefited from computational resources provided by the Australian Government through the National Computing Infrastructure (NCI). I'd especially like to thank Rodger Edberg (NCI) for supporting my work.

I would also like thank Jane Speechley (Charismatic Communications) for editing a draft of chapter 3.

While in the US I learned a great deal from a great number of generous people including: Ron Griffin, Ron Kaiser, Ralph Werbs (Texas A & M University), Zachary Donohew (University of Central Arkansas), David Zilberman, Michael Hanemann, Larry Karp, Leo Simon, Brian Wright, Tony Rossman and Daniel

Tregeagle (UC Berkeley), Jay Lund (UC Davis), Ellen Hanak (PPIC), Ron Anderson (LCRA), Jim Kassal (SWB), Jason Peltier (Westlands), Steve Hirsch and Craig Brackbill (MWDoSC), David Zetland (Leiden University), Steve Hvinden and others at (BoR Lower Colorado), Randi Field, Jon Traum and others at (BoR Mid-Pacific), Craig Trombly (DWR), Stu Townsley (US ACE) and Tim Washburn (SAFCA). I've also had email correspondence with Cars Hommes (University of Amsterdam) and Leigh Tesfatsion (Iowa State University).

I would like to thank Leo Simon (UC Berkeley) and Rachel Goodhue (UC Davis) for inviting me to the *77th International Atlantic Economic Society Conference* in Madrid 2014 and for providing feedback on chapter 5. While in Madrid I also had useful discussions with Alberto Garrido (Technical University of Madrid). I'd also like to thank Jesus Anton and Julien Hardelin for welcoming me at the OECD.

I received helpful feedback on chapter 8 from Akshay Shanker (ANU) and others at the *Crawford School Phd Conference 2014* and on chapter 7 from Paul Wyrwoll (ANU) and others at my final ANU seminar in 2015. I would also like to thank various attendees at the *Computing in Economics and Finance* conference in Vancouver 2013 and the *Agricultural Economics Annual Conference 2014* in Paris with which I had many thought provoking discussions.

Finally, I would like to thank my wife Kylie for being the most tolerant woman in the world and my daughter Elizabeth for being the least tolerant.

Abstract

This thesis is concerned with the allocation of water in regulated rivers: rivers controlled by large dams. Property right reforms undertaken in Australia and elsewhere have decentralised water allocation: both spatial (via market trade) and inter-temporal (via storage rights). While these reforms have proven successful, externalities and transaction costs always persist. This thesis asks: what is the preferred system of water rights — in terms of allocative efficiency — given that all proposals are in some way 'second-best'?

We test property right systems using a decentralised model of a regulated river, in which a large number of users make private trade and storage decisions. The model is unique in representing surface water rights in a multi-agent stochastic dynamic environment (a stochastic game). To solve the model we develop a novel computational method in the spirit of 'multi-agent systems', which combines reinforcement learning algorithms from computer science with learning concepts from game theory. Ultimately, this allows us to populate the model with near optimal selfish agents.

We present three applications of the model. First, we consider the design of water storage rights — where users hold private storage reserves in public reservoirs. In particular, we compare 'capacity sharing' (Dudley and Musgrave 1998) with alternative approaches. Second, we reconsider the issue of priority water rights — where certain users receive water allocations before others — in the context of storage. Third, we reconsider all of these issues, in the context of in-stream users: in particular a large environmental water holder. To maintain generality we specify broad ranges for the parameters of our model — based on statistics for the Australian Murray-Darling Basin (MDB) — and present the results of a large number of model runs. We compliment our results with discussion of the water institutions of Australia and the western US.

Contents

List of Figures

List of Tables

Model symbols

Indexes

t	Time period, $t \in [0, 1, 2, ..., \infty)$
i	Water user, $i \in [0, 1, 2, ..., n]$
h	User 'reliability' group, $h \in [low, high]$
j	River node, $j \in [1, 2, 3]$

Aggregate variables

I_{t+1}	Storage inflow
S_t	Storage volume
C_t	Climate state (i.e., dry or wet)
M_t	Season (i.e., summer or winter)
W_t	Storage release (withdrawal)
L_t	Storage evaporation losses
F_{jt}	River flow at node j
Z_t	Storage spills
E_t	River extraction
R_t	Return flow
Q_t	Water use
A_t	Water allocation
P_t	Spot market water price

User level variables

q_{it}	Water use
e_{it}	Productivity level
a_{it}	Water allocation
u_{it}	Payoff
s_{it}	Water account balance
w_{it}	Water account withdrawal
l_{it}	Water account loss deduction
x_{it}	Water account externality

Parameters

β	Discount rate
K	Storage capacity
ρ_I	Inflow autocorrelation
c_v	Inflow coefficient of variation
δ_0	Evaporation rate
α	Storage surface area parameter
δ_a, δ_b	Delivery loss parameters
$\theta_{h0}, ..., \theta_{h5}$	Irrigator yield function parameters
$\mathscr{A}_h$	Irrigator land area
ρ_e	User productivity shock autocorrelation
λ_i	User i inflow share
Λ_{high}	High reliability user group share of inflow
τ	Spot market transfer cost
ψ_h	Risk aversion parameter
ρ_C	C_t autocorrelation
$b_\$, b_j$	Environmental benefit function parameters
δ_R	Return flow rate
δ_{Ea}, δ_{Eb}	Extraction loss parameters

Functions

$\mathcal{L}_0(S_t, .)$	Evaporation loss function
$\mathcal{L}_1(F_{1t})$	River delivery loss, between node 1 and 2
$\mathcal{L}_2(F_{1t})$	River delivery loss, between node 2 and 3
$\mathcal{L}_E(F_{1t})$	Extraction loss function
$\mathcal{R}(E_t, .)$	Return flow function
$\pi_i(q_{it}, .)$	Consumptive user profit function
$d_i^{-1}(q_{it}, .)$	User i inverse demand function
$B(.)$	Environmental benefit function
ν_h	Utility function for user group h
$G_I(.)$	Unconditional cumulative distribution of inflow

Distributions

$U(a, b)$	Uniform over range (a, b)
$N(\mu, \sigma^2)$	Normal with mean μ and variance σ^2
$N_a^b(\mu, \sigma^2)$	Truncated normal: $N(\mu, \sigma^2)$ truncated to the range (a, b)
$\Gamma(k, \theta)$	Gamma with mean $k\theta$ and variance $k\theta^2$

Stochastic shocks

ϵ_{t+1}	Inflow / climate shocks, $\epsilon_{t+1} \sim \Gamma(k_I, \theta_I)$
η_{it}	User i productivity shock, $\eta_{it} \sim N(0, \sigma_\eta^2)$
ω_t	Seasonal inflow shock, $\omega_t \sim N_{\omega_a}^{\omega_b}(\mu_\omega, \sigma_\omega^2)$

Units of measurement

Volume

1 megalitre (ML)	= 10^6 litres (L)	= 0.81 acre-feet (AF)
1 gigalitre (GL)	= 1000 ML	
1000 gigalitres	= 10^5 ML	= 0.81 million acre-feet (MAF)

Area

1 hectare (HA)	= 10,000 meters^2 (m^2)	= 2.47 acres (ac)

Length

1 kilometer (km)	= 1000 meters (m)	= 0.62 miles (mi)
1 millimeter (mm)	= 0.001 m	= 0.039 inches (in)

Online resources

All of the code written for this thesis is available online via github.

github.com/nealbob/regrivermod

See appendix B for an introduction to the code.

Further information is available on my website

nealhughes.net

including additional model results, more detail on the computational techniques and related conference papers and publications.

Acronyms

ABARES Australian Bureau of Agricultural and Resource Economics and Sciences.

ABS Australian Bureau of Statistics.

ACE Agent Based Computational Economics. , 201

ASGD averaged stochastic gradient descent. , 192

AWA Australian Water Association.

BoR United States Bureau of Recalamation. , 39, 72

CBD Central Business District. , 2

CEWH Commonwealth Environmental Water Holder. , 139

CS capacity sharing.

CSIRO Commonwealth Scientific and Industrial Research Organisation.

CVP Central Valley Project. , 72

DWR Californian Department of Water Resources. , 72

ENSO El-Nino Southern Oscillation. , 21

EWH Environmental Water Holder. , 139

HL A two level 'high-low' priority right system. , 112

ID irrigation district. , 73

IE Institutional Economics. , 58

LCRA Lower Colorado River Authority. , 73

LP linear programming. , 37

MDB Murray-Darling Basin. , 1, 17

MDBA Murray-Darling Basin Authority.

MDBMC Murray Darling Basin Comission.

MDBMC Murray Darling Basin Ministerial Council.

MDP Markov decision process.

MPE Markov perfect equilibria. , 196

MWDoSC Metropolitan Water District of Southern California. , 73

NASA National Aeronautics and Space Administration.

NIE New Institutional Economics. , 63

NS no storage rights. , 83

NSW New South Wales. , 17

NWC National Water Comission.

OA open access storage. , 83

OE oblivious equilibria. , 197

OLS ordinary least squares. , 189

PPIC Public Policy Institue of California.

QLD Queensland.

RBF radial basis function. , 189

RL reinforcement learning. , 178

RS release sharing.

SA South Australia. , 17

SAFCA Sacramento Area Flood Control Agency.

SDP stochastic dynamic programming. , 37

SGD stochastic gradient descent. , 190

SWA spillable water account. , 81

SWP California State Water Project. , 72

TC tile coding.

UC University of California.

US United States.

USACE United States Army Corps of Engineers. , 31

USGS United States Geological Survey.

VIC Victoria. , 17

Chapter 1

Introduction

1.1 Motivation

1.1.1 Droughts, floods and climate change

Australia's 'Millennium drought' was the worst in recorded history (CSIRO 2010). The effect on agriculture was dramatic: between 2000 and 2010 the government paid $4.5 billion in drought assistance to farmers (ABARES 2012). Meanwhile, the water supply to cities was threatened, with most households on water restrictions for the best part of a decade.

At the height of the drought (2006 to 2009) inflows to the Murray River were near record lows three years running[1]. The major dams in the Murray-Darling Basin (MDB) fell to critical levels and water allocations to priority irrigators — normally immune from droughts — were cut significantly. The Murray River ceased to flow at the mouth for several years, threatening an environmental catastrophe (Taylor 2008).

The desperation of the times was reflected in the words of then Prime Minster John Howard:

> It's very serious, it's unprecedented in my lifetime ... we should all, literally and without any irony, pray for rain.
> (John Howard on ABC Television, 19th April 2007).

[1]Murray River inflows for 2006-07 were just 1040 GL, almost half the previous record low of 1920 GL set in 1914-15. Around 2500 GL were received in 2007-08 and 2100 GL in 2008-09. The long-term average is around 11,000 GL. (MDBA 2014*a*)

The drought ended with a series of flooding events. In 2010, the major dams of the southern MDB filled and flooding occurred in northern Victoria. These events were trivial in comparison with those of early 2011. By the 11th of January, three quarters of Queensland had been declared a flood disaster zone (Hurst 2011) — including the CBD of Brisbane. The next day the rain moved south to Victoria, causing "one of the biggest floods in the states history" (Premier Ted Baillieu, cited in Willingham 2011).

While droughts and floods are nothing new for Australians, these events were particularly severe. Subsequent research has confirmed that climate change contributed to the Millennium drought (CSIRO 2010). In fact, these recent events are indicative of a worldwide trend towards greater climate extremes (Coumou and Rahmstorf 2012, Seneviratne et al. 2012).

In addition to increased volatility, climate change is expected to induce large regional shifts in rainfall. In Australia, reductions in rainfall and stream-flow are expected in the already water stressed south-eastern regions (CSIRO 2008*b*, NASA 2013).

1.1.2 An era of dam building

During the 20th century Australia's main response to water scarcity was the construction of large dams[2]. Large dams served to smooth the supply of water from rivers, especially to provide water during droughts. This in turn facilitated large scale development: both major coastal cities and inland irrigation areas.

Dam construction peaked in the 1960s before halting in the 1980s (figure 1.1). A similar trend has been observed internationally (figure 1.2).

While dam construction has slowed in developed countries — largely due to a lack of good sites — a second boom in dam building is now underway in developing countries (e.g., China, Brazil and Africa) driven mostly by hydro-power (Verhoeven 2012).

[2]The International Commission On Large Dams (ICOLD) define large dams as those greater than 15 metres in height or 3GL in volume.

Figure 1.1: Australian storage capacity 1901- 2005

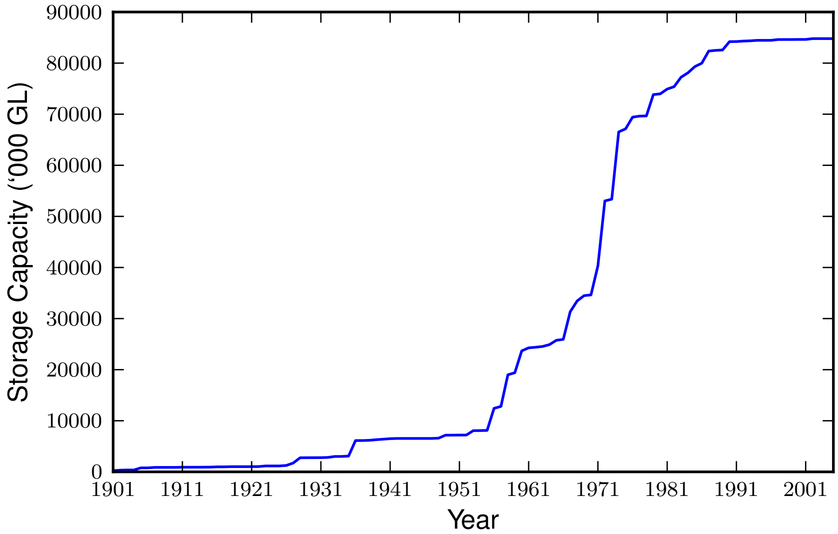

Source: ABS (2015)

Figure 1.2: Number of large dams placed into operation by decade (World)

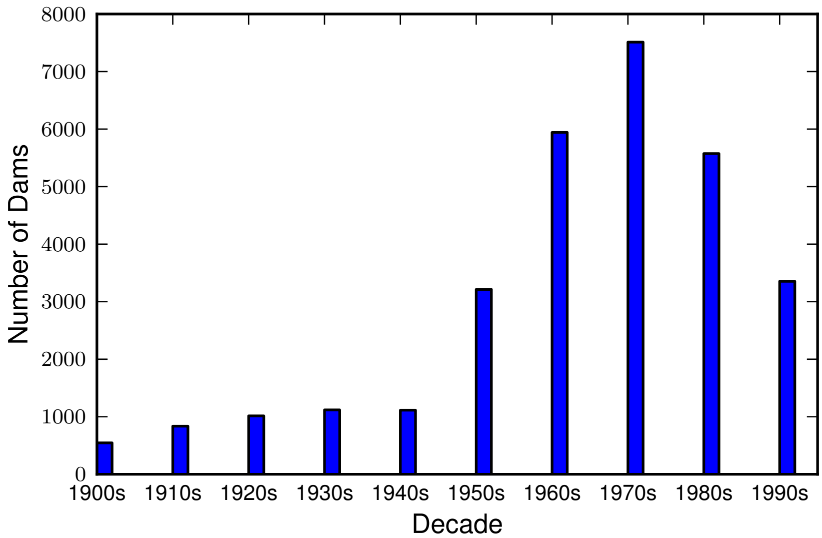

Source: ICOLD (2007)

This enthusiasm for dam building has long been criticised, both in Australia and elsewhere (Davidson 1969, McCully et al. 1996, Hirshleifer et al. 1969). Studies have shown that dam construction costs and times often well exceed estimates (Ansar et al. 2014) and benefits — particularly irrigation profits — can go unrealised (Davidson 1969, McCully et al. 1996)[3].

Further, we now have a better understanding of the many side effects of large dams (McCully et al. 1996), particularly their environmental consequences. Large dams alter river flow patterns: reducing average flows, decreasing variability and reversing seasonal trends. These changes, have had dramatic effects on river and wetland ecosystems (Poff and Zimmerman 2010).

The Colorado River in the western US is symbolic of the problem. Following the completion of the Hoover Dam in 1936, flows to the Colorado River Delta in Mexico essentially ceased. Once a thriving 7,800 km^2 estuary, the delta is now a small barren mud flat overrun with invasive species (Glenn et al. 1996).

1.1.3 Demand for environmental flows

In 2007 at the height of the drought, the Australian Government announced a $10 billion water policy package[4]. Central to the policy was a 'Basin Plan' (MDBA 2014*b*) to set new limits on extraction. Ultimately, the policy amounted to a large compensated transfer of water from irrigation farmers to the environment.

To say the plan has been controversial is an understatement. Farmer groups were extremely vocal in their opposition (Gale et al. 2014), environmentalists felt it did not reallocate enough water to the environment (Wentworth Group 2014) and economists argued there was too much emphasis on infrastructure upgrades (Grafton 2010).

In November 2014, the final plan — involving a reduction in extraction of around 20 per cent or 2750GL (MDBA 2014*e*) — was passed by parliament[5]. This water is held by a government agency: the Commonwealth Environmental Water Holder (CEWH) who uses it to achieve environmental goals.

While the Water Act was significant — both in its scale and the degree of control it gave the Commonwealth — it was indicative of a long-term policy trend towards

[3]In fact, the push for better evaluation of dam projects was central in the evolution of cost benefit analysis in economics (Griffin 2012).

[4]This budget was increased to $ 12.9 billion in 2008.

[5]As of 20th June 2014 1904 GL of this water has already been acquired through a combination of market purchase and infrastructure projects (MDBA 2014*e*).

environmental restoration[6]. Efforts to secure environmental flows in regulated rivers are now occurring in many parts of the world. In early 2014, the first of a series of 'pulse' flows was released into the Colorado River Delta[7].

1.1.4 Water property rights: trade and storage

Over time, water policy has gradually shifted from engineering and central planning towards economics and property right reforms. Economists prescription has long been water markets (Hirshleifer et al. 1969, Burness and Quirk 1979, Randall 1981). Market reforms have since occurred in many parts of the world, and nowhere with as much enthusiasm as Australia.

After decades of reform, Australia has established — in the southern MDB — perhaps the world's most efficient water market (Grafton et al. 2011*b*). While far from perfect, the market has proven successful in getting water to 'highest value uses' quickly with minimal involvement from government, especially during the recent drought (NWC 2011*b*).

However, trading only addresses one part of the water allocation problem: how we allocate our water across space. The other often neglected part, is how we manage our storages: how we allocate water across time.

Storage management involves difficult trade-offs between: maintaining supply 'reliability', minimising storage losses, providing environmental flows and minimising flood risk. Historically, our storages have been managed centrally according to simple rules, developed at a time when demands were relatively low and climate change was unheard of.

In hindsight, storage policy in the MDB was too liberal — storage reserves were too small — in the lead up to the drought. While the drought could not have been predicted, storage policy in the 90s was highly myopic. The standard practice was to 'over allocate': to allocate more water than could ever be used (MDBMC 1995). Given the 'use it or lose it' nature of water rights (Brennan 2008*a*) farmers consumed water until it had near zero value.

In response, there has been a trend toward decentralised management of storages, through the definition of water storage rights. These rights allow users to maintain private storage reserves in public reservoirs and make their own storage decisions.

[6]Previous policies included The Murray-Darling Basin Cap (the Cap), The Living Murray Initiative (TLM) and the National Water Initiative (NWI).

[7]This relatively small 130 GL release was the first step in a planned five year experiment, under an agreement between Mexico and the US.

While storage rights reform has lagged market reform (many regions introduced storage rights near the end of the drought) storage rights are now common place in the MDB (Hughes et al. 2013). Similar arrangements exist overseas. A form of storage right was even introduced for Hoover Dam in 2007[8].

1.2 The research question

Water is a uniquely complex commodity. Both the demand and supply of water are subject to non-rivalries, non-linear relationships and a high degree of climate related uncertainty. As such, defining exclusive property rights to water is difficult. In practice, externalities persist and trade remains subject to friction — regardless of how 'well' property rights are defined.

While markets may achieve gains over central planning, they'll always be some distance from optimal. In this type of second-best environment (in the sense of Lipsey and Lancaster 1956) it's difficult to make predictions about specific water reforms.

The question posed by this thesis is: what is the preferred approach to water rights in regulated rivers? That is: which property rights system maximises allocative efficiency, given that each proposal imposes different externalities and has varying reliance on market transactions?

The benchmark here is aggregate social welfare: including the private benefits of extractive water use (i.e., irrigation profits) and the public benefits of in-stream use (i.e., the environmental benefits of river flows).

Another way of thinking about the research question is in terms of the limits of markets. That is, setting aside some of the practical constraints, how far can property right reforms be pushed, how many decisions can we decentralise, how many externalities can we internalize, before we reach the limits imposed by the physical characteristics of water?

Before continuing we should define some boundaries. Firstly, we are concerned exclusively with regulated rivers — rivers where the flow is controlled by one or more large dams — as opposed to unregulated rivers or groundwater.

Secondly, we are concerned with the allocation of water over existing infrastructure. Here we take both the supply side (i.e., dams) and demand side (i.e., irrigation) infrastructure as fixed, putting aside any asset expansion or rationalisation.

[8] These storage rights are known as 'Individually Created Surplus' (ICS) and are held by lower Colorado water contractors (irrigation districts). ICS was introduced under the 2007 'interim shortage guidelines' (see Hughes 2013).

Finally, while this analysis is intended to be general our focus is largely on the Australian MDB. We also consider the western US, particularly the Colorado River, central California and southern Texas.

1.3 Method

Our method is essentially the computational experiment. Our basic philosophy is to represent as much of the complexity as is practical and/or necessary in a computational model, then impose alternative policy scenarios (i.e., property rights systems) and observe which performs best.

As we want to represent market imperfections — externalities and transaction costs — we require a truly decentralised model. That is, a model where the agents (i.e., water users) each respond optimally (i.e., selfishly) to the incentives imposed by a given property rights system.

1.3.1 The model

This thesis is concerned with the allocation of water within an abstract regulated river system (figure 1.3).

Our 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.

The water allocation problem facing the hypothetical social planner involves three components: storage (the inter-temporal allocation), extraction (the trade-off between consumptive and in-stream use) and use allocation (the allocation across consumptive users).

In the decentralised model, each agent makes private water storage and trade decisions subject to a set of water right (i.e., water accounting) rules. The users attempt to maximise their own welfare, taking into account their water right endowment, the accounting rules and the behaviour of the other agents. The role of government is limited to setting and implementing the water accounting rules.

To maintain generality, we specify distributions for the parameters of our model — reflective of Australian rivers — and solve our model for a large number of parameter values. This allows us to test how the performance of water right systems depends on the nature of the river.

Figure 1.3: An abstract regulated river system

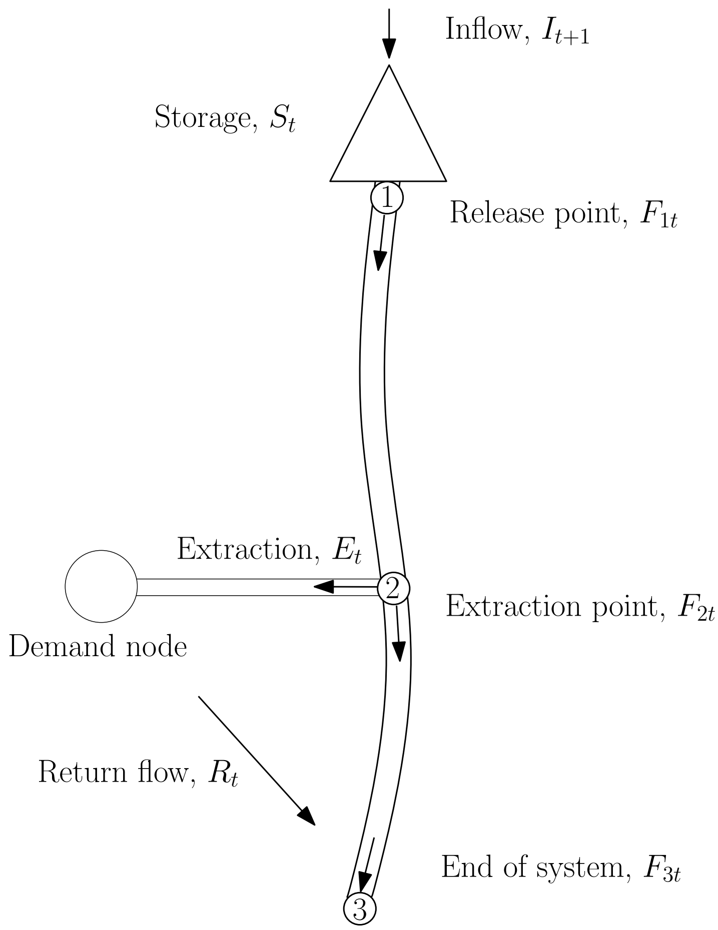

1.4 Contributions

1.4.1 Reconsidering water property rights in the context of storage

Water property rights have received a great deal of attention from economists. However, few studies have considered water rights in the context of storage. Traditionally, storage has been the domain of engineers, and economists have focused on use allocation and extraction.

This thesis offers a reconsideration of water property rights in the context of storage. To this end, we draw inspiration from the work of Norman Dudley (Dudley et al. 1971, Dudley and Burt 1973, Dudley 1988*a*, Dudley and Musgrave 1988, Dudley and Hearn 1993, Dudley et al. 1998), particularly his capacity sharing proposal — a system of water property rights where users hold shares in both reservoir capacity and inflows. We also draw much inspiration from the work of the late Donna Brennan (Brennan and Scoccimarro 1999, Brennan 2006; 2008*a*)[9].

To date the work of Dudley, Brennan and others in this area has relied on centralised (i.e. social planner) models, precluding formal analysis of water rights. No existing study has modelled second-best (i.e., non-exclusive) water property rights in the context of storage. The obvious reason is that a decentralised stochastic dynamic model with externalities is hard to solve. While not attempted for surface water, problems of this type have been attempted for groundwater and other natural resources.

Armed with such a model, we extend the literature in a number of directions. Firstly, we consider water storage rights, in particular we compare capacity sharing with alternative systems. Secondly, we consider the issue of priority water rights — where certain users receive water before others — in the context of storage and storage rights. Thirdly, we reconsider both of these issues, in the context of in-stream demands, particularly large environmental water holders.

1.4.2 Solving large stochastic games by reinforcement learning

Formally, our decentralised problem is a stochastic game (Shapley 1953). In stochastic games each user faces a Markov Decision Process (MDP) where the

[9] This thesis is also closely related to the work of Steve Beare (Beare et al. 1998; 2006, Beare 2010), Chi Truong (Truong et al. 2010, Truong and Drynan 2013), and Freebairn and Quiggin (2006). It also extends the authors own research on water storage rights in the MDB (Hughes and Goesch 2009*b*, Hughes 2010, Hughes et al. 2013).

payoff and transition functions depend on the behaviour of the other agents. Such problems arise frequently in natural resources and industrial organisation.

This thesis introduces a new method for solving stochastic games with large numbers of players. In the spirit of 'multi-agent systems' (Fudenberg and Levine 2007) our method combines reinforcement learning algorithms from computer science with concepts of learning in games from economics. Ultimately, these methods allows us to populate the model with 'intelligent' (i.e., near optimal) agents.

In reinforcement learning algorithms, agents 'learn' by observing the outcomes of their actions (i.e., optimisation by simulation). Similar to dynamic programming, these methods exploit the Bellman (1952) principle. Despite their obvious applicability, such methods have seen virtually no use by economists (Fudenberg and Levine 2007, Tesfatsion and Judd 2006)[10].

Our method extends the reinforcement learning literature in two directions. First, we modify the approach of fitted Q iteration with tile coding (Timmer and Riedmiller 2007), to make it suitable for economic problems. Second, we develop a multi-agent algorithm suited to large complex stochastic games, where the existence and uniqueness of Nash type equilibria is difficult to establish.

Our approach provides a middle ground between dynamic programming methods used in macroeconomics (like that of Krusell and Smith 1998) and methods used in Agent-based Computational Economics (ACE) (Tesfatsion and Judd 2006) — both in terms of the types of models it can be applied to and the degree of rationality assumed for the agents.

1.5 Outline

This thesis can be divided into three main parts: background (chapters 2, 3 and 4), results (chapters 5, 6 and 7) and methods (chapter 8 and appendices A and B).

1.5.1 Background

Chapter 2 details the physical characteristics of water as a commodity. Here we outline the hydrology of regulated rivers and consider consumptive (e.g., irrigation) and in-stream (e.g., environmental) demands. Throughout we draw on examples from the Australian MDB and western US.

[10]These computer science methods are not to be confused with the 'foresight free' methods sometimes applied in repeated games (Erev and Roth 1998).

Chapter 3 formalises the water allocation problem as a planner's stochastic dynamic optimisation problem. Here we also consider the role of government as central planner in water allocation. Finally, the chapter presents numerical results, contrasting optimal and myopic storage policy.

Chapter 4 introduces water property rights markets in the context of our model. The chapter also considers the literature on institutions in natural resources allocation. Here we attempt to reconcile our property rights approach, with the broader view of water institutions as complex 'polycentric' systems. Finally, the chapter compares Australian and US water rights and markets.

1.5.2 Results

Chapter 5 considers water storage rights. Here we introduce the decentralised version of the model, and use it to compare various approaches to storage rights including capacity sharing, spill forfeit rules, open access storage and no storage access. These approaches vary in how they represent the storage capacity constraint and evaporation losses.

Chapter 6 considers prioritisation in water rights in the context of storage. Here we consider traditional water rights — where a planner makes storage decisions — and capacity sharing each with and without priority water flow rights. The chapter considers two motivations for priority rights from the literature: minimisation of trade requirements and risk aversion.

In chapter 7 we add a large Environmental Water Holder (EWH) to our decentralised model. Our EWH is treated identically to all other users under the water rights framework; except their withdrawals (releases) remain in the river. Here we consider whether the presence of an EWH — with payoffs defined over river flows — changes any of our conclusions regarding the design of water storage or flow rights.

1.5.3 Methods

Chapter 8 details the computational methods used to solve our model. The chapter summarises the literature on stochastic games, including equilibrium concepts and solution methods. We then detail our methods and how they relate to existing approaches in computer science and economics.

Appendix A details our parameter assumptions. Here we present a statistical analysis of the major storages in the MDB — which informs the supply side parameters — and an econometric analysis of ABARES irrigation survey data — which informs the demand side parameters.

Appendix B provides further details on the implementation of the model, including an introduction to the code.

Appendices C, D and E provide further details on chapters 5, 6 and 7 respectively.

Chapter 2

Water as a commodity: supply and demand

2.1 Introduction

Water becomes a commodity when available in a form, at a location and at a time suitable for human use, or when it can cost effectively be transformed, transported or stored.

The majority of human water supply is extracted from rivers[1]. Water in rivers can be stored behind dams and transported by gravity through watercourses and channels. The extreme spatial and temporal variability of these resources, means that water is — at various locations and times — a scarce commodity.

This thesis focuses on regulated rivers — rivers controlled by dams — for three reasons. First, regulated rivers provide the majority of water supply in developed countries. Second, the economic analysis of groundwater is, in many respects, further developed than that of surface water. Third, the stocks and flows of water within regulated rivers can be accurately measured and controlled (at least relative to other hydrological systems).

Water demand can be classified consumptive or non-consumptive. Consumptive use transforms or transports water, such that it exits the effective supply. For example, water applied to crops may be 'lost' via evapotranspiration, while water used in households becomes waste to be disposed of (or recycled).

[1] In Australia in 2010-11 92 per cent of 'distributed' water supply was extracted from rivers, with 6 per cent coming from groundwater and 2 per cent from desalination (ABS 2012*b*). In 2011-12 Australian agricultural water use was 83 per cent surface water, 17 per cent groundwater (ABS 2012*c*). In 2005, total US water use was 80 per cent surface water, 20 per cent groundwater (USGS 2009).

In Australia and the US, the major consumptive user is agriculture, accounting for around half of all use in both countries (ABS 2012*b*, USGS 2009). Domestic consumption accounts for 25 per cent in Australia and 38 per cent in the US (ABS 2012*b*, USGS 2009).

Non-consumptive 'use' is where the flow of water generates value, for example: hydro-power generation or environmental benefits. In practice, the distinction between consumptive and non-consumptive use is approximate. For example, some water applied to crops may return to the river (return flows), while environmental flows may increase evaporation losses. As Hirshleifer et al. (1969; p. 66) note: "consumptive use of water should be thought of as a matter of degree".

In this thesis, the terms non-consumptive use and in-stream use are used interchangeably. Water use is taken to mean consumptive use unless otherwise stated.

This chapter details the physical properties of water as a commodity. We begin with water supply, detailing the inflow, storage and delivery of water within regulated rivers. We then turn to water demand, focusing on irrigated agriculture and environmental flows.

This chapter seeks to illustrate the complexity of water, providing a brief introduction to the hydrology, agronomy and ecology of regulated rivers. In the next chapter, we present a simplified model which captures the main economically relevant features.

While the focus is general, we include examples from rivers in the Australian MDB and the western US. Throughout the chapter, we present statistics on rivers in the MDB. For further detail on these statistics see appendix A.

2.2 Water supply

A stylised regulated river is shown in figure 2.1. Most regulated rivers can be considered generalisations of this framework, just with multiple storages, multiple inflow sources and multiple demand nodes.

Figure 2.1: A regulated river system

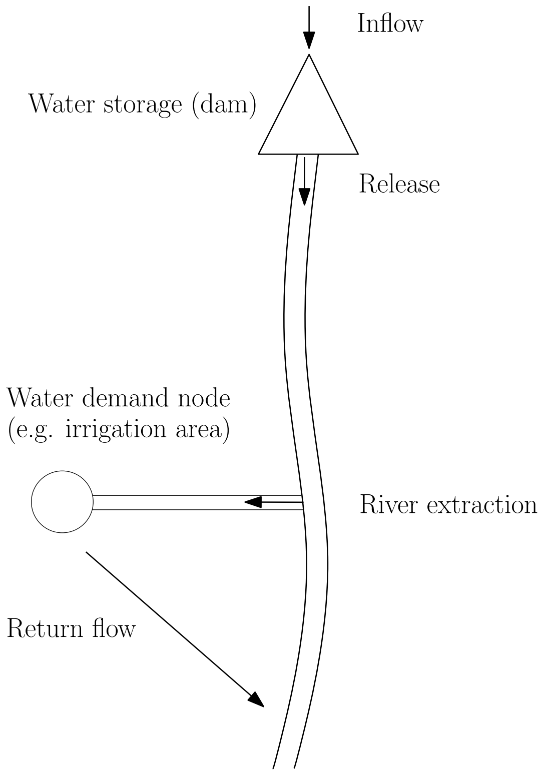

A regulated river can be described by three components, inflow, storage and delivery. River flows are a product of precipitation in an above catchment area. A dam is a barrier which impounds flow in an artificial lake (reservoir). Water stored in a reservoir can be delivered to users via rivers, channels and pipes.

Before continuing. we briefly consider two economically significant and heavily regulated rivers, the Murray in Australia and the Colorado in the US.

Figure 2.2: The Colorado River Basin (left) and Murray-Darling Basin (right)

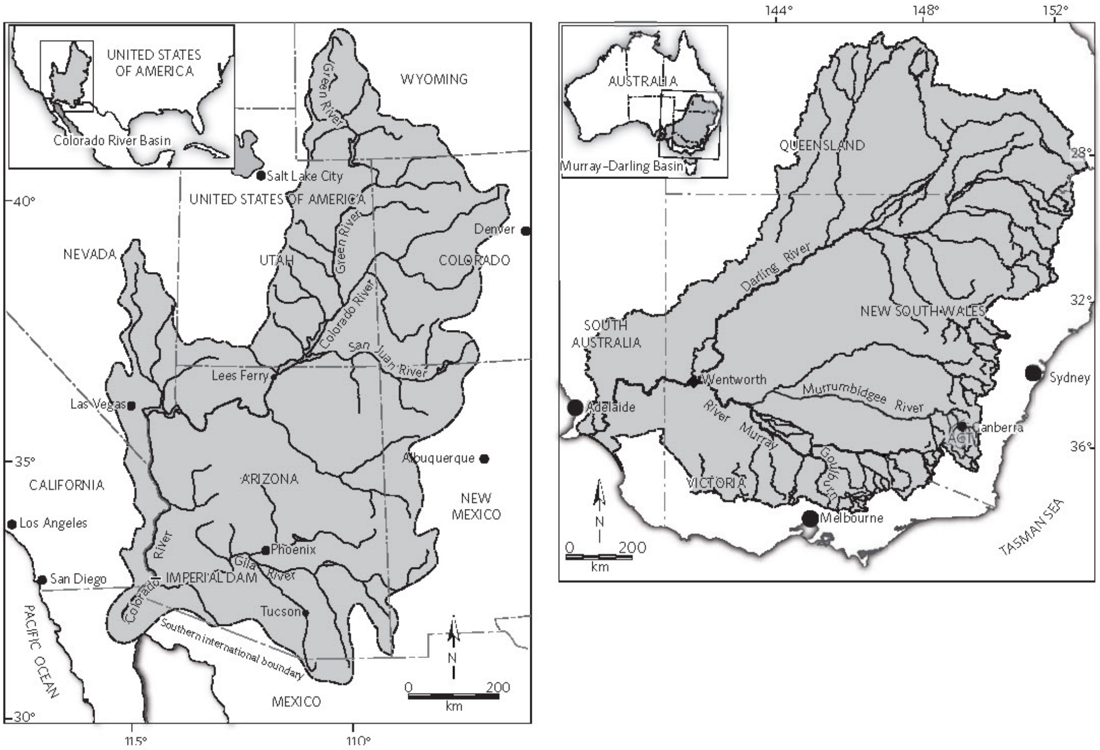

Source: Grafton et al. (2012)

Table 2.1: Comparison of the Murray-Darling and Colorado River Basins

	Murray-Darling Basin	Colorado River Basin
Basin land area (km^2)	1,061,469	618,000
Proportion of national land	0.14	0.08
Average Annual Flow (GL)	13,000	20,000
Storage capacity (GL)	24,500	74,000
Elevation change (meters)	2,120	3,104
Share of national farm income	0.39	0.15
Agriculture share of water use	0.83	0.78
Number of states	5	7
Population (million)	2	12.7
Share of national population	0.10	0.04

Source: MWDoSC

2.2.1 The Murray

The Murray-Darling Basin (MDB) drains an area of over 1,000,000 km^2 in South-Eastern Australia (figure 2.2). The Murray River extends 2,375 km, forming the border between New South Wales (NSW) and Victoria (VIC), before entering South Australia (SA) and draining to the sea near Adelaide (figure 2.3). Along the way, the Murray receives tributary flow from the Darling, Goulburn and Murrumbidgee rivers, among others.

Figure 2.3: The Murray River

Source: MDBA (2012*b*)

The vast majority of Murray inflow is runoff from the Australian Alps in the southeastern corner of the basin. For the size of the catchment, Murray inflows are very low and variable. During the recent 'Millenium drought' (2000-2010) MDB inflows were 40 per cent below the long run average.

Inflow is captured by the Hume Dam (3000 GL, completed 1936), Dartmouth Dam (3850 GL, completed 1979) and the Snowy Mountains Hydroelectric Scheme (constructed between 1949-1974) — a network of dams, power stations and tunnels which divert water (that previously flowed east into the Snowy River) into the Murrumbidgee and Murray Rivers. Total storage capacity in the MDB is 24,500 GL — around twice mean annual inflow.

Murray water is used primarily for irrigation, but also supplies the city of Adelaide. Water is delivered from Hume Dam to irrigation areas in NSW, VIC and SA, via a series of weirs and locks, before entering channel networks. A downstream off-river storage, Lake Victoria, regulates flow into SA.

2.2.2 The Colorado

The Colorado River Basin drains an area of 637,000 km^2 in the southwestern US and northwest Mexico (figure 2.2). The Colorado river travels 2,334 km from the Rocky Mountains in Colorado, through Utah, Arizona, Nevada and California before crossing into Mexico (figure 2.4).

Figure 2.4: The Colorado River

The majority of inflow comes from snowmelt off the Rocky Mountains. By US standards inflows are highly variable. The Colorado has experienced prolonged drought since 2000 with inflows in the order of 30 per cent below the long run average between 2000-2010.

The major storages on the Colorado include Lake Mead (Hoover Dam, 35,000 GL, completed 1936) and Lake Powell (Glen Canyon Dam, 32,000 GL, completed 1966).

The Colorado has an enormous storage capacity of 74,000 GL — around four times the average annual inflow. A number of smaller downstream dams including Parker Dam (completed 1938) provide short-term river regulation.

The Colorado is subject to high extraction, with typically no flow reaching the sea. The majority of the water is for irrigation, much of it in the south of the basin. The Colorado also supplies major cities including Los Angeles, Las Vegas and Phoenix. In many cases, Colorado water is transported by channels vast distances before use. For example, the Colorado River Aqueduct transports water 400 km into Los Angeles and the Central Arizona Project covers 541 km into Arizona.

2.2.3 Inflow

Precipitation and runoff

Inflow is ultimately a product of precipitation: rain and snow. Rain and snowmelt either evaporate, recharge groundwater or generate surface run-off. The formation of runoff depends on evaporation rates, soil type, antecedent soil moisture, vegetation and land slope, among other factors.

Commonly a high proportion of run-off is generated in mountainous regions, with high precipitation, low evaporation and steep slopes. Run-off from rain is often concentrated in storm events. At high altitudes snowpack can form in winter and melt during spring.

Runoff is affected by water 'interception' activities in the catchment, such as farm dams and plantation forestry (see for example Van Dijk et al. 2006). Rivers can also gain or lose flow through connection with groundwater[2].

A convex relationship is typically observed between rainfall and runoff: the proportion of rainfall that forms runoff increases as rainfall increases (Davie 2008). As a result, variation in streamflow generally exceeds that of rainfall (figure 2.5).

Inflow as a random variable

For our purposes inflow can be treated as a random variable. Inflow distributions can be estimated from historical data (obtained from stream gauges or derived from rainfall or tree-ring records). While inflow distributions are catchment specific, they

[2] The CSIRO (2008*a*) found that MDB rivers tend to be gaining in the upper parts of a catchment and losing in the lower parts

Figure 2.5: Annual MDB rainfall and natural Murray River flow 1900-2009

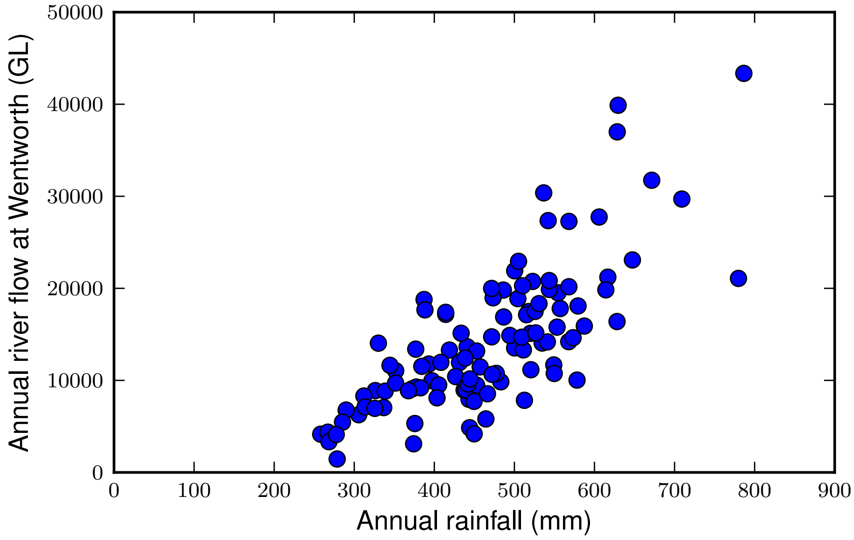

Source: MDBA (2012*b*)

Figure 2.6: Modelled natural Murray River flow (at Hume Dam) 1900-2009

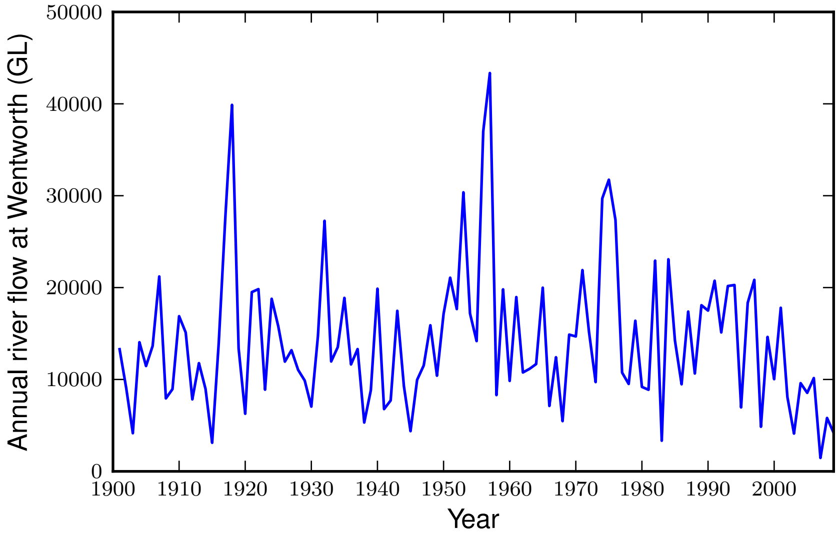

Source: MDBA (2012*b*)

share a number of characteristics: high variation, positive skewness[3], seasonality and positive autocorrelation (McMahon et al. 2007).

Flow variation can be viewed along various time scales: short run (i.e., daily or hourly), within year (i.e., seasonal), annual (figure 2.6) and multiple year cycles. Seasonality varies by region: inflows are winter dominant in the southern MDB and summer dominant in the north. In the Colorado inflow peaks in spring due to snowmelt.

Of particular interest, are the frequency and severity of inflow extremes, both floods — brief high inflow events — and droughts — extended periods of low inflow.

The severity of drought depends importantly positive autocorrelation in inflow. Positive autocorrelation is due in part to ocean temperatures cycles, including the El-Nino Southern Oscillation (ENSO). These multi (2-15) year cycles have persistent effects on the spatial distribution of rainfall. In Australia, the El-Nino phase of ENSO is associated with dryer than average conditions.

Australian rivers are particularly variable, having 1 to 3 times less average flow (per unit catchment area) and around double the flow variation (coefficient of variation) compared with the rest of the world (McMahon et al. 2007).

Figure 2.7: Annual inflow coefficient of variation, Australian rivers

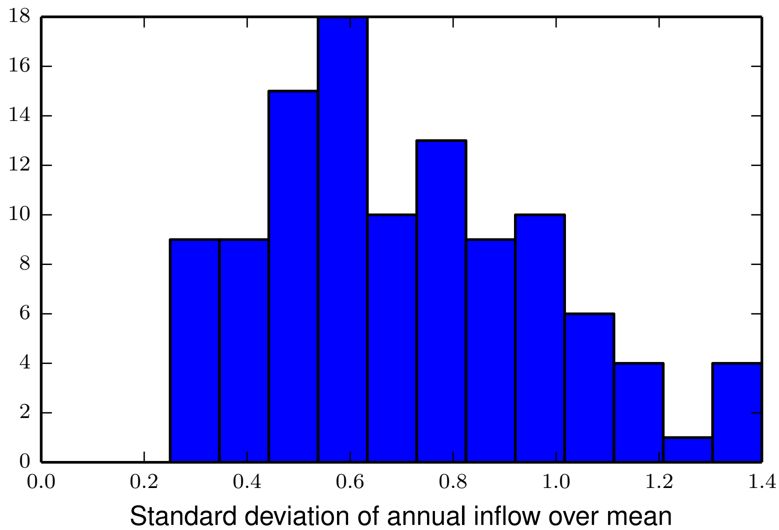

Source: Peel et al. (2010)

[3]Stream-flow distributions commonly have heavy right tails (Katz et al. 2002)

Table 2.2: Estimated annual autocorrelation, selected MDB rivers

River	Location	$\hat{\rho}_I$
Murray	Yarrawonga Weir	0.23
Murrumbidgee	Burrinjuck Dam	0.28
Goulburn	Lake Eilden	0.28
Namoi	Keepit Dam	0.23
Ballone	St George	0.22

Source: MDBA (2012*b*)

Climate change

Historically, inflow distributions were assumed stationary, however long run shifts are now expected due to climate change.

Climate change is expected to exacerbate existing spatial and temporal rainfall variation: dry regions are likely to get dryer, wetter regions likely to get wetter, and the severity of extremes is likely to increase (see Kundzewicz et al. 2008).

While estimates are subject to much uncertainty, reductions are expected in both the MDB and Colorado Basins. The CSIRO (2008*b*) predicted a reduction in average water supply in the MDB of 11 per cent by 2030. A reduction in average flows of 14 per per cent has been predicted for the Colorado by 2040(Christensen et al. 2004).

These climate change predictions are now being supported by growing evidence. The CSIRO (2010) concluded that climate change was a significant contributing factor to the Millennium drought. Coumou and Rahmstorf (2012) summarise evidence linking recent increases in the frequency of extreme weather events (such as floods and heatwaves) to climate change.

2.2.4 Storage

Large dams

The primary purpose of most large dams is to smooth variability in the supply of water (figure 2.8). Other purposes include hydro-power generation, river navigation and flood mitigation. The most common primary purpose is water supply for irrigation (48 per cent), followed by hydro-power (17 per cent), domestic water supply (13 per cent) and flood mitigation (10 per cent) (ICOLD 2007).

Figure 2.8: Dartmouth Dam monthly inflows and releases 1990 to 2012

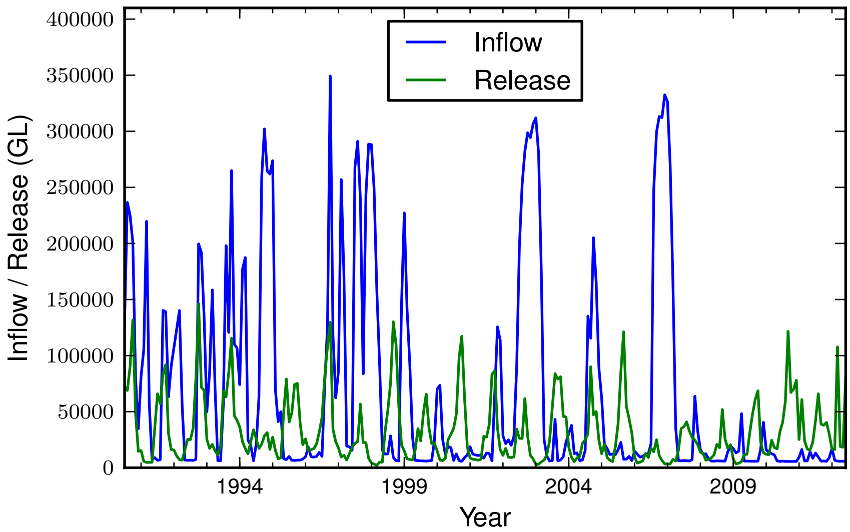

Source: MDBA (2015)

Figure 2.9: Histogram of storage capacity, MDB storages

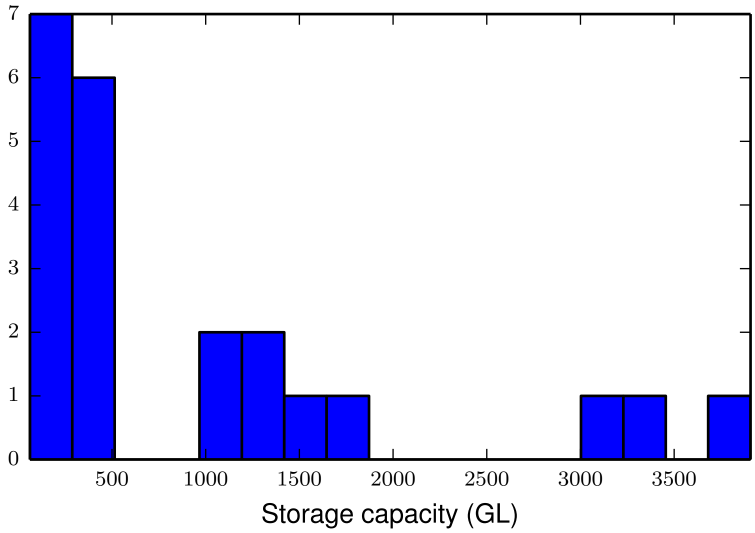

Source: NWC (2011*a*)

Storage capacity, spills and losses

Storage capacity is a function of dam height and the geometry of the landscape. Ideally, dams are located in a steep valley to maximise volume to surface area. While often treated as a constant, capacity can decline slightly over time due to sedimentation. Reservoir volume below the dam release outlet — generally a small proportion of total volume — is known as dead storage as it is inaccessible without pumping.

When at full capacity inflows spill uncontrolled from a dam and flow downstream. Spills are seen as losses to consumptive users, since they tend to occur when demand is low. However, spills can have in-stream benefits or costs, such as environmental effects or flood damage.

The frequency of spills depends on the mean and variance of inflow relative to capacity. In the Colorado River, spills are infrequent (Hoover dam's only spill occurred in 1983). In contrast, spills are an almost yearly occurrence in the Northern MDB (due to small dams and volatile inflows).

Water is lost from storage via evaporation and seepage. In Australia, storage losses are almost entirely due to evaporation (Gippel 2006). Storage loss depends on climate as well as storage geometry. Evaporation losses are assumed proportional to surface area (Lund 2006). The relationship between volume and surface area is generally concave (Lund 2006).

Figure 2.10: Annual mean inflow over capacity, MDB storages

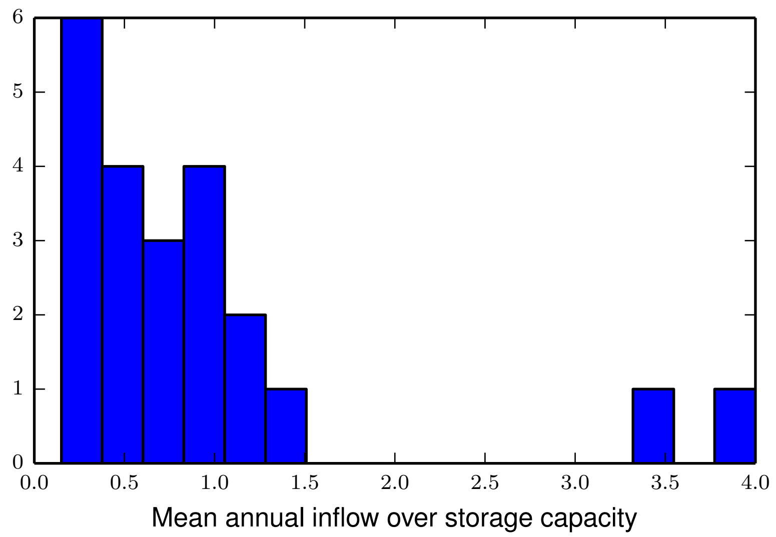

Source: NWC (2011*a*)

Figure 2.11: Surface area over capacity, MDB storages

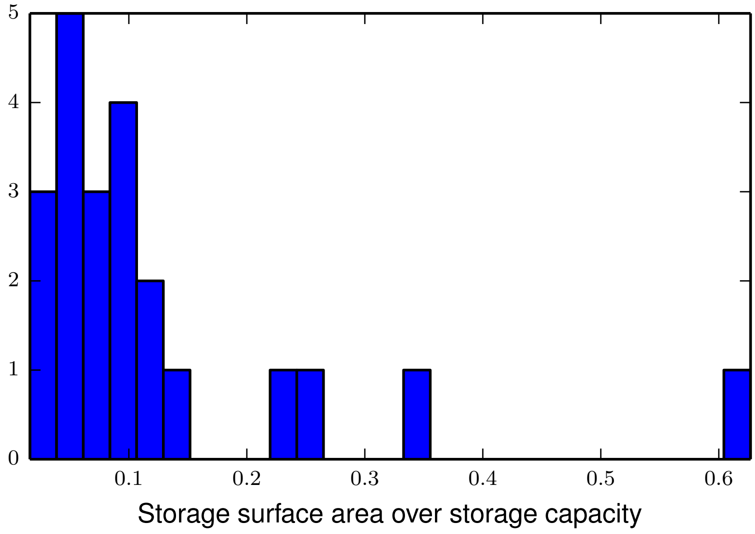

Source: ANCOLD (2013)

Off-river storage

In some cases dam releases or spills, may be stored in downstream locations. One option is to divert water into off-river reservoirs: like the Imperial Dam on the Colorado, San Luis reservoir in central California or Lake Victoria on the Murray. These reservoirs tend to be smaller, subject to higher evaporation losses and are generally used for short term regulation rather than inter-year reserves.

Shallow on-farm storages (ring tanks) are another option. Farm ring tanks are common in some Northern MDB catchments where dams are small and relatively inefficient.

Another option is 'groundwater banking': the artificial recharge of groundwater aquifers. A key advantage of groundwater banking is a reduction in evaporation losses and spills. Groundwater banking is limited in Australia, but is undertaken on a large scale in some Western US regions (examples include Kern County California and the 'Arizona Water Bank').

The possibility of large scale groundwater banking in Australia has been investigated by the BRS (2007) and Ross (2012). Challenges include the costs of recharge (either by injection bore or infiltration basin), seepage losses, deterioration in water quality and the extraction costs.

2.2.5 Delivery

Delivery by river

On the Murray major extraction begins at Yarrawonga Weir 180 km down river from the Hume Dam (65 km straight line distance and four days flow time). Flow speed depends on the river's elevation profile: the Murray is a slow meandering river in comparison with the Colorado.

Rivers can both gain and lose flow. Gains can come from tributary streams, groundwater or return flows. Typically, delivery involves a net loss due to evaporation and seepage. Delivery losses depend on distance, climate, the nature of the channel and the rate of flow.

In general, high marginal losses are expected at low flow rates, while the river bank is relatively unsaturated, and at high flow rates, when over bank flows become significant. In-between there may be flow rates where marginal losses are relatively low (NWC 2010).

A number of US studies estimate losses in arid zone rivers with intermittent flows (Cataldo et al. 2010). This literature supports the model of Lane (1983), involving 100 per cent loss below a flow threshold and a constant rate of loss thereafter.

Estimates of losses in Australian rivers vary across studies and are subject to significant measurement error (Gippel 2006). Losses over regulated river sections can be relatively low. On the Murray, losses average less than 4 per cent of annual flow (or 167 GL) between Hume Dam and Yarrawonga Weir (Gippel 2006). Losses tend to be higher and more variable in the less regulated lower reaches of the Murray (Gippel 2006).

A common assumption in hydrological models is a piecewise linear relationship between flow and loss (MDBA 2011): 100 per cent loss below a low flow threshold (i.e., a fixed loss), a range of zero marginal loss and a constant marginal loss above a high flow (i.e., over bank flow) threshold. Losses between Hume Dam and Yarrawonga Weir are assumed almost entirely fixed (MDBA 2011).

Delivery by canal

Most irrigation areas are supplied by canals. For example the Mulwala Canal delivers water from the Yarrawonga Weir to the NSW Murray irrigation area. Typically, an initial main canal gives way to a network of smaller channels within the irrigation area.

Delivery losses per unit distance are generally higher in canals than in rivers. In the southern MDB, losses are around 20-30 per cent (Hume 2008). Losses include evaporation and seepage and outfalls (unused water leaving the end of the canal system). Significant losses can be incurred in filling the canals at the beginning of the season.

Losses to evaporation and seepage depend on climate, canal surface (earthen or lined), flow rate and canal geometry. Griffin (2006) considers the relationship between flow rates and losses in optimally designed irrigation canals, finding losses to be concave functions of flow.

Hume (2008) presents annual data on total extraction and losses in southern MDB irrigation areas. At this level there is support for a linear relationship between flow and loss, with a significant fixed component (figure 2.12).

Figure 2.12: Annual losses, NSW Murray irrigation area, 1995 to 2009

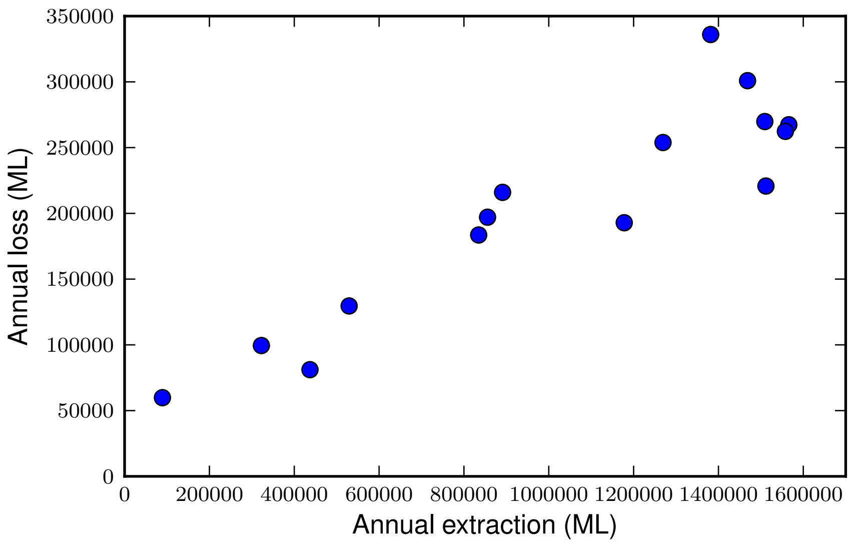

Source: Hume (2008)

2.3 Water demand

2.3.1 Irrigated agriculture

Irrigation farms generate 30 per cent of Australia's agricultural production (13 billion in 2010-11) from less than 0.5 per cent of its agricultural land. The MDB accounts for around half of Australian irrigation — \$5.9 billion in 2010-11 (ABS 2012*a*).

In dry environments moisture is the limiting factor in crop growth. Soil moisture is a combination of natural moisture and irrigation water. Plant growth increases with soil moisture until a point of maximum yield.

Irrigated agriculture can be divided into three industries: horticulture, broadacre and dairy. Most broadacre crops involve a summer growing season. Perennial horticulture (e.g., fruit trees) involves life cycles in the order of 30 years with annual harvest seasons. With perennials water stress can affect current and future yields — under extreme water stress trees may be destroyed.

Water is one of many inputs in the farm production function. However, there are generally limited options for substitution in the short term. Between seasons broadacre farms can alter crop types and areas planted. In the long term, all farms can invest in irrigation technology, change industries or revert to dryland farming.

Water input demand has been estimated from a variety of techniques including: optimisation models, econometric studies and field experiments. Scheierling et al. (2006) provides a meta analysis of 24 US studies, Hughes (2011) summarises the Australian studies.

The literature confirms a number of expected results: water demand is inversely related to price (Scheierling et al. 2006); inversely related to rainfall (Brennan 2006, Hughes 2011); more price elastic in broadacre than in horticulture (Scheierling et al. 2006, Bell et al. 2007, Hughes 2011); and more price elastic in the long run (Scheierling et al. 2006).

Figure 2.13 shows farm yield functions — mapping water use per unit land to profit per unit land — for broadacre and horticulture farms are estimated using ABARES farm survey data for the southern MDB (see appendix A).

Figure 2.13: Estimated profit-water functions

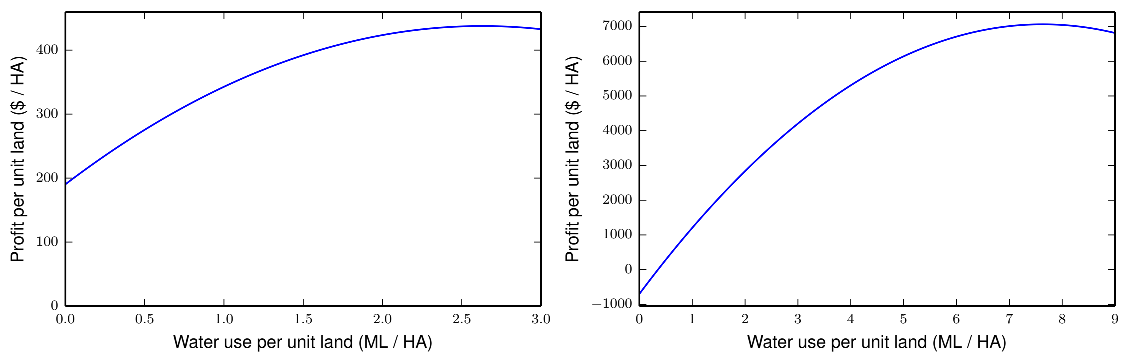

(a) Low reliability (broadacre)

(b) High reliability (horticulture)

2.3.2 Environmental water demand

The consequences of regulation and extraction

Storage and extraction alter natural river flows (figure 2.14): lowering average flows, reducing variability and changing seasonal flow patterns. These alterations have significant ecological consequences (Poff and Zimmerman 2010). The MDBA (2010) provide a detailed account of the environmental effects of regulation and extraction in the MDB.

Figure 2.14: Modelled Murray river flows at the SA border

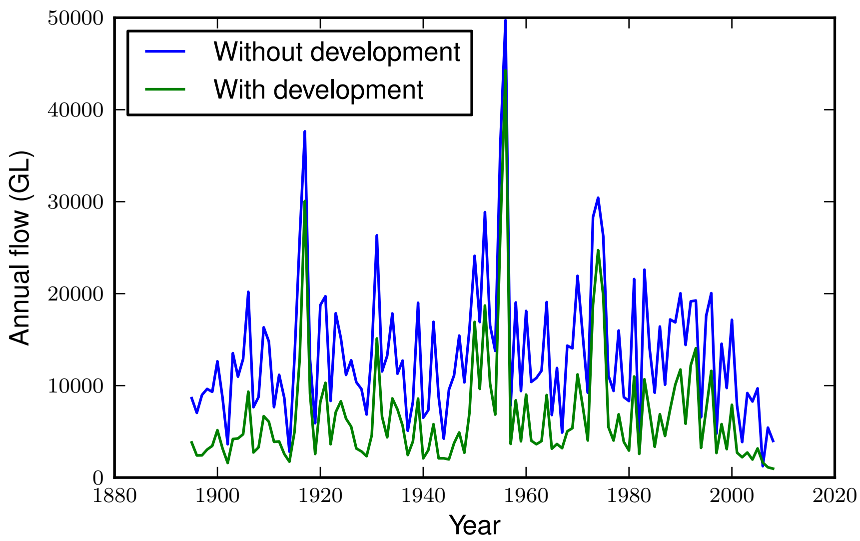

Source: MDBA (2012*b*)

Reductions in flow are felt most in the lower reaches of a river, often damaging wetlands at the mouth. Since the completion of the Hoover Dam flows to the Colorado River Delta have been negligible. The delta is now 5-10 per cent of its former size. In the MDB, the health of the Murray Mouth-Coorong-Lower Lakes wetlands has been affected by reduced flows. In 2009-10 lake levels declined below sea level increasing salinity and exposing acid sulphate soils (Kingsford et al. 2011).

Regulation reduces the frequency of high flow events which inundate wetlands. On the Murray the frequency of large floods into the Barmah-Millewa Forrest has reduced significantly. Between 2000 and 2010 there were no large flood events, damaging trees (particularly river red gums) and limiting bird breeding events (MDBA 2012*a*).

On the Murray, regulation has reversed seasonal flow regimes from winter to summer dominant. These changes have important consequences for plants and animals with annual breeding cycles. For example, unseasonal (summer) flooding into the Barmah-Millewa Forrest can damage wetland trees (Chong and Ladson 2003).

Environmental flows

Environmental flows are storage releases designed to replicate aspects of natural flow regimes. Generally, environmental objectives (e.g., increasing mean river flow and/or increasing flow variability) are in conflict with consumptive demands. Large flow events may also be in conflict with flood mitigation goals.

Estimating the benefits of environmental flows involves both a scientific problem: understanding ecosystem responses to changes in river flows and an economic (non-market valuation) problem: estimating the social value of ecosystem improvements. While the risk of ecological damage is known to increase with the magnitude of flow alteration, estimating precise relationships is difficult (Poff and Zimmerman 2010). In addition, there remains much uncertainty over non-market valuations of river ecosystem values (CSIRO 2012, PC 2010).

Environmental flows are often based on a comparison of the current and the pre-development (i.e., natural) flow distributions, leading to a set of flow targets. For example, minimum average flows or a maximum time between flood events. The MDBA (2010) defined many potential flow targets for the Murray River, some examples are shown in table 2.3.

Table 2.3: Example environmental flow targets

Site	Flow target	Percentage of years required
Barmah-Millewa	25,000 ML per day for 6 weeks	40-50
	50,000 ML per day for 3 weeks	25-30
Lower lakes, Coorong	3 year average > 2000 GL per year	95
	3 year average > 1000 GL per year	100

Source: MDBA (2010)

2.3.3 Flood mitigation

Flood mitigation involves capturing large inflow events in storage to prevent flood damage in downstream areas. Flood mitigation ‘pre-releases’ can be made when storage volumes are close to capacity or the probability of a storm is high. These releases free storage space, so that larger storm peaks can be captured, rather than spilling uncontrolled. As with environmental flows, flood mitigation is in conflict with water supply: generally pre-releases are equivalent to spills for consumptive users.

Flood mitigation is typically secondary to water supply in the MDB (MDBA 2014*c*). Flood mitigation is a more significant issue in wet and highly populated regions. Flood mitigation is a higher priority for many US dams, particularly those constructed by the United States Army Corps of Engineers (USACE). A prime example is Folsom Dam, located above the capital of California, Sacramento.

2.4 Conclusion

The majority of water supply in developed countries is extracted from rivers. As a product of precipitation, river flows are naturally variable. However, variation in streamflow generally exceeds that of precipitation, due to non-linear rainfall-runoff relationships. By international standards Australian rivers are particularly variable.

Of most concern to users are droughts: extended periods of low rainfall, high temperatures and high evaporation, leading to both low river flows and high water demands. A key driver of droughts are multi-year climate cycles such as the ENSO.

The relationship between water supply and willingness to pay (the demand curve) is typically non-linear (convex). During droughts, the scarcity value of water can increase dramatically as shortages are imposed on 'higher value' users: such as horticulture or households. In contrast, the value of water can quickly become zero (or negative) during floods.

These variations in the value of water create demand for storage. The most common form being large dams. Large dams exist to smooth the supply of water, greatly increasing the size of development (agricultural or urban) that is possible.

As a form of commodity storage, reservoirs involve a number of unique characteristics. Storage capacity is fixed (in the short run), constrained by height of the dam and the geometry of the landscape. With variable inflows this constraint can easily bind resulting in spills. Further, storage surface area and evaporation losses are non-linear (concave) functions of volume.

One problem with large dams is that they dramatically alter natural river flows. Large dams are associated with reductions in the mean and variance of flow and seasonal changes. These changes weaken the connections between rivers and floodplains and have a range of environmental consequences.

Recognition of environmental damage has created demand for environmental flows: storage releases designed to replicate aspects of natural flow regimes. Environmental flows are examples of 'in-stream' water demands: where the flow of water within a river generates public benefits. Another key example is flood mitigation.

The central conclusion from this chapter is that water is a complex commodity. Both the supply and demand for water are characterised by: exogenous variability, non-linear relationships and non-rivalries. The remainder of this thesis is concerned with how different types of institutions handle this complexity.

Chapter 3

The water allocation problem: storage, extraction and use

3.1 Introduction

The water allocation problem in a regulated river can be divided into three parts: storage, extraction and use.

The storage decision is an inter-temporal problem: use water now or store for future periods. During droughts — with both low inflows and high water demands — storage reserves are vital for water users exposed to high losses from shortages.

Storage policy involves a yield-reliability trade-off. Higher storage reserves reduce the variability of supply, but increase losses from spills and evaporation. The optimal policy will depend on storage constraints, inflow probabilities and the nature of water demands.

The extraction decision involves a trade-off between consumptive and in-stream use. In the absence of in-stream demands, the trivial solution is to extract all flow. In practice, in-stream demands are significant and — particularly in the case of environmental flows — are increasing in importance over time.

The final decision is how to allocate extraction across individual users. This can be viewed as a spatial problem; for example allocating water to farmers across an irrigation area. The ideal allocation will depend on the marginal benefits of use and water delivery constraints.

In this chapter, the water allocation problem in a regulated river is defined as a social planners stochastic dynamic optimisation problem. While the social planner's problem is introduced largely as an instructive device, water allocation is often the domain of central government agencies. In this chapter, we briefly consider the role of central planners in water allocation, before turning to property rights and markets in chapter 4.

At the end of the chapter, we develop a parametric version of our planner's problem and present some illustrative numerical results. This model forms the basis of the decentralised models developed in chapters 5, 6 and 7.

3.2 The planners problem

This thesis is concerned with the allocation of water within an abstract regulated river system (figure 7.1). The system involves a single storage supplying water to many consumptive users located at a single demand node.

Figure 3.1: An abstract regulated river system

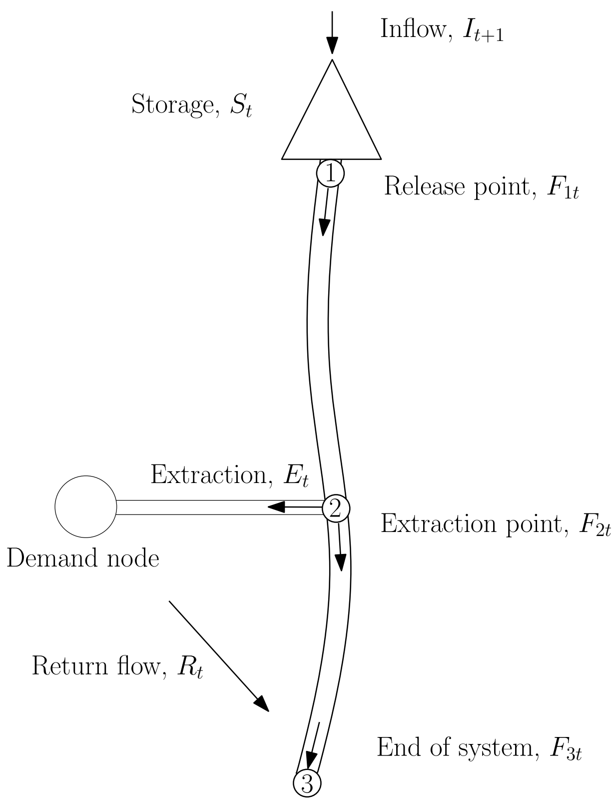

The model is in discrete time with an infinite horizon. For now, the time periods $t \in [0, 1, 2, ..., \infty)$ are arbitrary. There are n water users $i \in [1, 2, ..., n]$ and three river flow nodes $j \in [1, 2, 3]$ (see figure 7.1).

The model is presented in general form below; a parametric example is presented in section 3.5.

3.2.1 Water supply

Each period the storage receives inflow I_{t+1}. Inflows are variable and subject to positive autocorrelation and seasonality. Formally I_{t+1} is a continuous random variable with conditional density $g(I_{t+1}|C_{t+1}, M_{t+1})$. C_{t+1} is a proxy for the climate state of the region, such that $\text{Cov}(I_{t+1}, C_{t+1}) > 0$. C_t follows a Markov process with $\text{Cov}(C_{t+1}, C_t) > 0$. M_{t+1} indicates the season.

The storage level S_t evolves according to

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

$$0 \leq W_t \leq S_t$$

Here K is the storage capacity, W_t is the storage release (i.e., withdrawal) during period t and L_t is the storage (evaporation) loss. Storage losses are a non-decreasing and concave function of S_t

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

River flow volumes F_{jt} during period t are defined for each node j. Flow below the storage F_{1t} equals releases plus storage spills Z_{t+1}:

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

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

River flows downstream of the extraction point F_{2t} are

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

$$0 \leq E_t \leq F_{1t} - \mathcal{L}_1(F_{1t})$$

Here E_t is river extraction and $\mathcal{L}_1$ is the delivery loss function (non-decreasing in F_{1t}, $\mathcal{L}_1(F_{1t}) \in [0, F_{1t}]$). End of system flows F_{3t} are

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

where $\mathcal{R}$ is the return flow function (non-decreasing in E_t, $\mathcal{R}(E_t, M_t) \in [0, E_t]$) and $\mathcal{L}_2$ is another delivery loss function.

3.2.2 Water demand

There are $i = 1$ to n water users (e.g., irrigators) each with water benefit (e.g., profit) function $\pi_i(q_{it}, C_t)$, where q_{it} is water 'used' by user i during period t. Here, the climate state C_t acts as a proxy for natural moisture availability (e.g., irrigation area rainfall). π_i is concave in q_{it} and C_t. Water use q_{it} and climate C_t are substitutes. The inverse water demand function is

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

Total water use Q_t is constrained by the volume of water delivered to the demand node: extraction E_t less delivery losses $\mathcal{L}_E(E_t)$

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

$$Q_t \leq E_t - \mathcal{L}_E(E_t)$$

In-stream demands are represented by an objective function over river flows $B(F_{1t}, F_{2t}, F_{3t}, .)$. For now in-stream demands remain general and could refer to environmental benefits, flood mitigation or hydro-power generation etc. We consider environmental objectives in chapter 7.

3.2.3 The objective

The planner's problem is to maximise the expected discounted benefits from consumptive and in-stream use, subject to the water supply constraints. Assuming consumptive and in-stream benefits are measured in the same units (e.g., dollars) the planner's problem is

$$\max_{\{q_{it}, W_t, E_t\}_{t=0}^{\infty}} E\left\{\sum_{t=0}^{\infty} \beta^t \left(\sum_{i=1}^{n} \pi_i(q_{it}, C_t) + \sum_{j=1}^{3} B_j(F_{jt}, C_t)\right)\right\}$$

subject to the constraints detailed above.

This is a stochastic dynamic optimisation problem with $n + 2$ policy variables (W_t, E_t, q_{it}) — the storage, extraction and use allocation decisions — and three state variables (S_t, C_t, M_t). In section 3.5 we present a numerical solution the problem.

3.3 Modelling literature

Engineers and economists have long developed computer models of surface water allocation problems. While attempts at 'integrated' models exist, most studies can be classified as economic or engineering (Dudley and Hearn 1993).

Economic studies emphasise the demand side, often combining detailed water use models (i.e., farm production models), with simple supply side assumptions (exogenous water supply or a single storage). Engineering studies emphasise the supply side, often focusing on complex river systems (with many storages), while using simple water use objectives.

Generally, engineering studies focus on optimal water allocation rules, while economic studies focus on the welfare differences between arbitrary and optimal scenarios: typically to support the adoption of market reforms.

3.3.1 Engineering literature

For a detailed review of the engineering literature, see Yeh (1985), Wurbs (1993) or Rani and Moreira (2010).

Engineering studies focus mostly on the storage decision (reservoir operation). A key problem is the management of multiple reservoir systems: how to distribute storage reserves to minimise total evaporation losses and spills[1].

Another focus is 'multi-purpose reservoirs': those with consumptive and in-stream demands. However, in-stream demands are often limited to hydro-power and flood mitigation. Often, in-stream and consumptive objectives are specified in different units and pragmatic 'multi-objective' methods are used[2].

Numerical techniques have evolved with improvements in computing power. Early studies used linear programming (LP), dynamic programming and stochastic dynamic programming (SDP) optimisation methods (Yeh 1985).

Early optimisation models were severely limited in the amount of detail that could be incorporated — motivating a reliance on simulation (Yeh 1985, Wurbs 1993). More recently there has been a trend towards combining simulation and optimisation, through the use of genetic algorithms and machine learning methods (Rani and Moreira 2010).

[1] A good example here is the work of Perera and Codner (1996; 1988). With multiple storages, the optimal approach depends on their spatial arrangement (e.g., in series or parallel, see Lund 1996; 2006).

[2] For example, maximising one objective at a time with the other is specified as a constraint or using various weighting schemes (see Wurbs 1993).

3.3.2 Economic literature

Traditionally, economists have focused on use allocation, particularly on estimating the gains from water trade. Early examples include Flinn and Guise (1970) and Vaux and Howitt (1984). Appels et al. (2005) summarises the Australian studies.

These static models derive the optimal allocation of water across users and space, via LP or quadratic programming in the vain of Takayama and Judge (1964). More recent examples include models of the MDB (Grafton and Jiang 2011, Hafi et al. 2009, Adamson et al. 2007) and the CALVIN model of California (Draper et al. 2003).

While economists have long applied SDP to natural resources such as fisheries, forestry and groundwater (Kennedy 1981) they have rarely considered surface water problems. The work of Norman Dudley is a key exception (Dudley et al. 1971, Dudley and Burt 1973, Dudley 1988*a*, Dudley and Hearn 1993).

For example, Dudley (1988*a*) presents a planner's SDP model of the Namoi region. The model assumes a single reservoir with stochastic inflow, supplying water to cotton farmers. Water demand is based on a detailed farm production model. The model assumes a quarterly time scale and a three year planning horizon.

A key focus of Dudley (1988*a*) was the total area of irrigation land in the context of new developments. Dudley (1988*a*) showed how larger irrigation areas increase the variability of supply and eventually decrease mean profits[3].

More recent examples of dynamic surface water models include Beare et al. (1998), Brennan (2008*a*), Howitt et al. (2002), Iglesias et al. (2007) and Hughes and Goesch (2009*b*). Brennan (2008*a*) focuses on the Goulburn region, Hughes and Goesch (2009*b*) on the Murrumbidgee and Iglesias et al. (2007) on southern Spain.

A common finding, is that central reservoir policies are overly aggressive: storage reserves are too small (Brennan 2008*a*, Iglesias et al. 2007, Hughes and Goesch 2009*b*). These models all show that optimal storage can significantly reduce the costs of droughts, whilst modestly increasing average welfare.

Dudley (1988*b*), Hughes and Goesch (2009*b*) and Brennan (2008*a*) also consider storage rights. These studies are considered further in chapter 5. Economic studies which model environmental demands are reviewed in chapter 7.

[3] In this thesis, the focus is on established systems, where the irrigation area is fixed (or if anything decreasing).

3.4 Central planning in water

Natural resource allocation involves complex systems of institutions, combining government, private and community decision making. Water allocation in the Australian MDB and the western US are prime examples (see chapter 4).

In this context, the role for government extends to both planning and providing institutions to support decentralised allocation. In this section, we consider only the government's role as water planner. We return to these broader issues in chapter 4.

3.4.1 Some historical context

Central planning has a long history in water allocation. Anthropologists and historians often link the development of water infrastructure with the formation of government bureaucracy. A prominent example being Wittfogel (1957)'s theory of 'hydraulic empires'.

In the twentieth century, government control of water can be closely linked with the development of large dams. In the US, government agencies funding the construction of large dams — such as the Bureau of Reclamation BoR — typically maintain a role in the allocation of their waters.

Australia has a long history of state government control of water, which can be traced back to the 1884 Victorian Royal Commission on Water Supply led by Alfred Deakin. Deakin visited a number of countries and was apparently horrified by the complex and litigious approach to water allocation in the US (California in particular). Upon returning, he advocated state government control of water. Subsequently, the Victorian Water Act (1886) gave ownership and control of water to the Victorian Government. Similar legislation followed in NSW and SA (for a detailed account see Davis 1967).

3.4.2 The water bureaucracy

In practice, water planning involves complex bureaucracies with many organisations and sets of rules. For detail on the MDB see Challen (2000), for the western US see Wurbs (2013) or Libecap (2011*b*).

A water bureaucracy typically involves a hierarchy of organisations and rules, leading to a series of nested water allocation decisions. These progress from large spatial and temporal units (e.g., long-run sharing of water resources between states/nations) down to finer decisions (e.g., daily allocation of water to individual farmers)[4].

Each decision involves a combination of rules (section 3.4.3) and discretion. As with policy rules in other domains (see Taylor 1993), rules are either incompletely defined, or — as is often the case in water — are broken during periods of scarcity (Hughes 2014).

3.4.3 Water allocation rules

Some common water allocation rules are defined below in the context of planner's problem (assuming no storage or delivery losses for simplicity).

Storage rules

The simplest storage rule is to release all water up to a predefined limit

$$W_t = \min\{S_t, \bar{Q}\}$$
$$\bar{Q} \leq K$$

In the engineering literature, this is known as Standard Operating Policy (SOP) and $\bar{Q}$ as the target demand (Lund 1996). Often $\bar{Q}$ reflects the maximum feasible use volume: where marginal benefits reach zero or delivery constraints bind.

More conservative 'hedging' rules accumulate reserves when $S < \bar{Q}$ (Lund 1996).

This type of hedging rule broadly reflects storage policy in Victorian MDB systems, while NSW is closer to the SOP (Hughes and Goesch 2009*b*).

Flood mitigation

Flood mitigation rules aim to keep downstream flows below critical thresholds. A typical rule involves a storage trigger $\bar{S}$ and threshold $\bar{F}$

$$W_t = \begin{cases} \min\{S_t, \bar{Q}\} & \text{if } S_t < \bar{S} \\ \min\{S_t - \bar{S}, \bar{F}\} & \text{if } S_t \geq \bar{S} \end{cases}$$
$$\bar{S} < K$$

[4] We consider the hierarchical structure of water institutions further in chapter 4.

Figure 3.2: Stylised storage rules

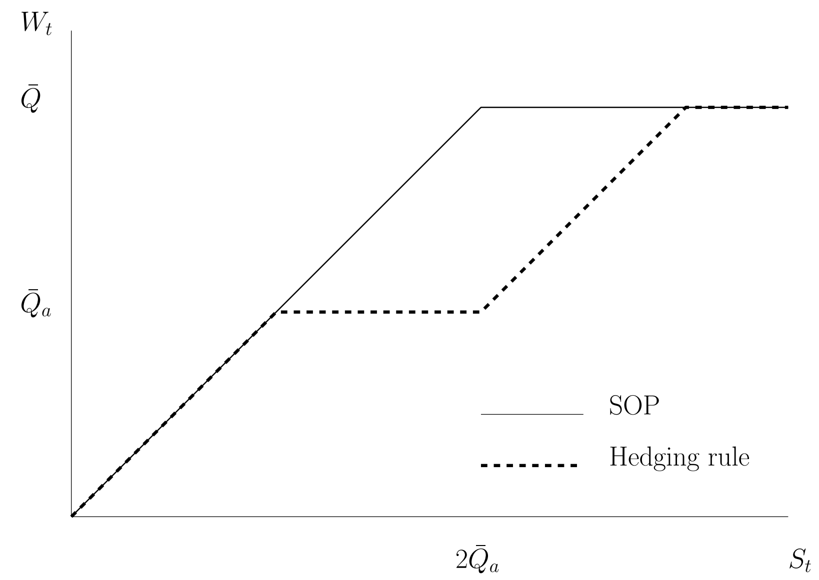

This type of flood mitigation rule is strictly enforced in US reservoirs, where $S_t \in [0, \bar{S}]$ is known as conservation storage and $S_t \in [\bar{S}, K]$ as flood control storage (Wurbs 2013).

Environmental flows

Environmental flows (F_t^e) are storage releases which can not be extracted

$$W_t = F_t^e + \min\{S_t - F_t^e, \bar{Q}\}$$
$$E_t = W_t - F_t^e$$

Two common environmental rules are 'transparent' and 'translucent' flows, respectively

$$F_t^e = \min\{I_t, \underline{\mathrm{F}}\}$$
$$F_t^e = \omega I_t$$

where $\underline{\mathrm{F}} > 0$ and $\omega \in (0, 1)$ are constants.

In the MDB 'rules based' environmental flows are commonly specified in water sharing plans[5].

[5] While not an environmental flow rule as such, a use or extraction cap (i.e., a lower $\bar{Q}$) can be used to achieve environmental goals. On its own, a cap will indirectly increase river flows by increasing average storage levels and therefore storage spills and flood mitigation releases.

Use allocation: proportional

The simplest use rule is to assign each user a proportional share λ_i of extraction

$$q_{it} = \lambda_i E_t$$
$$0 < \lambda_i < 1, \sum_{i=1}^{n} \lambda_i = 1$$

In irrigation areas, initial shares are often based on farm land area.

Use allocation: priority

More commonly, users are often divided into priority classes with water allocated to the 'senior' classes before 'junior'.

With two classes, the set of users $\mathcal{U} = \{i | i \in 0, 1, 2, ..., n\}$ is partitioned into two subsets, $\mathcal{U}_a$ and $\mathcal{U}_b$

$$q_{it} = \begin{cases} \min\{\lambda_i E_t, \lambda_i \bar{Q}\} & \text{if } i \in \mathcal{U}_a \\ \lambda_i(\max\{E_t - \sum_{\mathcal{U}_a} \lambda_i \bar{Q}, 0\}) & \text{if } i \in \mathcal{U}_b \end{cases}$$

This approach — with a small number of priority classes — is standard both in Australia and the western US (chapter 4). Typically, the classes reflect different user types: for example, households versus agriculture or perennial versus annual cropping farms.

Use allocation: priority ordering

Here the users $i = [0, 1, 2, ..., n]$ are ranked in order of priority (with $i = 0$ highest priority) then:

$$q_{it} = \min\{\lambda_i \bar{Q}, \max\{E_t - \sum_{j<i} \lambda_j \bar{Q}, 0\}\}$$

Priority ordering is a key component of the original 'Prior appropriation' approach to water rights developed in the western US (section 4.5).

3.4.4 The limitations of planning

The limitations of government control of natural resources are well established (see for example Libecap 2009, Ostrom 2010*b*, Grafton 2000). The historical failures of water policy in Australia and the US have also been well documented (Hirshleifer et al. 1969, Davidson 1969). Below we discuss the limitations of planning in the context of regulated rivers.

Planner's information problem

In section 3.2 we implicitly assumed the planner has complete information on both the supply of water (including the probability distribution over inflows) and the demand (the preferences of users).

In practice, governments have incomplete information on user preferences. The use allocation decision requires knowledge of individual preferences, which will depend on farm technology, climate and input and output prices among other factors. While the storage decision only requires aggregate preferences, information problems remain (Dudley and Musgrave 1988, Brennan 2008*a*, Hughes and Goesch 2009*b*).

Hierarchies are clearly an adaption to these information problems. Hierarchies allow high-level decisions to be made by central agencies, while finer decisions can be made by regional organisations with less spacial and social distance to users[6]. Policy rules are another adaption. While clearly not optimal, rules reduce the decision-making burden and provide some predictability for users.

In reality, inflow probabilities will also be unknown: there will be ambiguity (especially given climate change). Arguably, there is less asymmetry here: governments and users face the same ambiguity since inflow information is easily transferable. In this thesis, we assume inflow probabilities are known, and consider decision making under ambiguity outside our scope.

Government incentive problems

In section 3.2 we assumed the planner's goal was to maximise social welfare. In practice, government objectives can be politically biased (i.e., subject to rent seeking).

[6] This information based explanation for decentralisation in government is conventional wisdom in the economic literature on federalism (see for example Oates 1999). We discuss these institutional hierarchies in more detail in chapter 4.

Many authors have documented biases in water policy both in the US (Griffin 2012, Hirshleifer et al. 1969) and Australia (Watson 2005, Davidson 1969). Two common biases are the provision of water to agriculture at subsidised costs, and an over reliance on new infrastructure. Recent history provides much evidence of this second bias, including over investment in irrigation infrastructure in the MDB (Watson 2007) and desalination plants in cities (Brennan 2008*b*).

One response is to establish organisations with a degree of decision-making (and in some cases financial) independence. There are many recent Australian examples: the MDBA, environmental water managers (e.g., the Commonwealth Environmental Water Holder), and even storage managers (e.g., the Northern Victorian Resource Manager). However, recent experience — particularly the Murray-Darling Basin Plan process (Quiggin 2012) — suggests that it is extremely difficult to isolate these organisations from political realities.

Finally, corruption can arise: where government officials alter water allocation policy for personal gain. Some prime examples have emerged in Victoria in recent times (Victorian Ombudsman 2011).

Government policy uncertainty

Information and incentive problems explain why governments adopt sub optimal policies. A related problem is that governments generally can not commit to, or communicate, a complete water allocation policy ex ante.

In practice, users face investment decisions that depend on future water allocations. Where government policy is known, future allocations depend only on climate variation. In reality, government policy is an additional source of uncertainty, complicating planing and discouraging investment.

Policy uncertainty was a major focus of Dudley (1988*b*), who described it as a two-way information problem: planners don't know the preferences of users and users don't know the policies of planners.

> Reduced levels of economic efficiency are to be expected ... because of the interdependent nature of supply and demand probabilities, and the difficulty of deriving, summarising and communicating them between reservoir and farm managers. (Dudley 1988*b*; pp. 647)

The problem can be also be viewed as one of credibility. Governments tend to depart from water allocation rules frequently. These departures damage the credibility of future government policy commitments.

Further, the relationship between policy rules and user allocations can be complex, so that even with fixed rules, it can be difficult to calculate probability distributions over allocations. For example, allocations to users depend on environmental flow and flood control rules. These rules can be complex, even depending on external factors such as ecosystem or flood risk indicators.

Hughes et al. (2013) demonstrate the problem by showing large differences in the allocations made to users for equivalent levels of water supply over the period 1995 to 2010 in the MDB (figure 3.3).

Figure 3.3: Water allocation versus availability in the Murrumbidgee 1995 to 2010

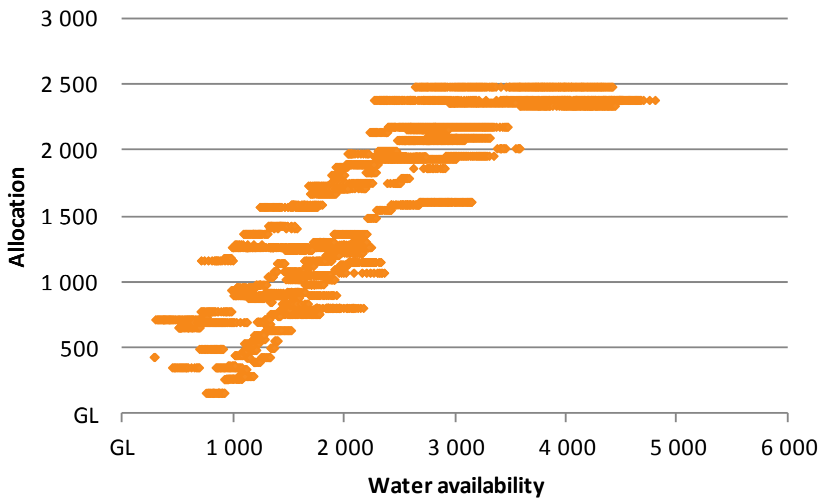

Source: Hughes et al. (2013)

Finally, the problem can be worsened by multiple priority classes, since the effect of rule changes tends to be concentrated on junior classes (chapter 6).

3.5 A parametric model

Below is a parametric version of the model, for now excluding in-stream demands and assuming an annual timescale.

3.5.1 Functional form

Inflow

Annual inflows are drawn from a gamma distribution, with first order autocorrelation

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

$$0 < \rho_I < 1$$

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

Here the climate variable C_{t+1} is just the lagged inflow I_t.

AR(1) log-normal and AR(1) gamma are common assumptions for annual inflows. In a study of 1221 international rivers, McMahon et al. (2007) find evidence that gamma fits inflow data better than log-normal.

Storage losses

A standard evaporation loss function is adopted following Lund (2006)

$$\mathcal{L}_0(S_t - W_t) = \delta_0.\alpha S_t^{2/3}$$

$$\delta_0, \alpha > 0$$

Here δ_0 represents the evaporation rate (per unit surface area) and $\alpha S_t^{2/3}$ the surface area.

Delivery losses

Without in-stream demands we can condense all delivery losses between the storage and the point of extraction into one function L_1

$$\mathcal{L}_1(F_{1t}) = \delta_a + \delta_b F_{1t}$$

$$0 < \delta_b < 1$$

$$0 > \delta_a$$

Note that without in-stream demands, we can ignore river flows and return flows downstream of the extraction point.

Consumptive water demand

Two types of user are defined: high-reliability (e.g., horticulture) and low reliability (e.g., broadacre). The set of users $\mathcal{U} = \{i \in 0, 1, ..., n\}$ is partitioned into $\mathcal{U}_{low} =$

$\{i \in 1, 2, ..., n_{low}\}$, $\mathcal{U}_{high} = \mathcal{U}_{low}^{\complement}$. The index $h \in (high, low)$ indicates membership to these sets.

For each class, a quadratic relationship is assumed between profit per land area, water use per land area $\tilde{q}_{it} = q_{it}/\mathscr{A}_h$ and climate conditions (in this case proxied by $\tilde{I}_t = I_t/E[I_t]$).

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

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 and e_{it} is a user specific productivity shock following an AR(1) process

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

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

$$0 < \rho_e < 1$$

3.5.2 The problem

With the above assumptions the problem condenses to:

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

Subject to:

$$S_{t+1} = \min\{S_t - W_t - \delta_0\alpha S_t^{2/3} + I_{t+1},\ K\}$$

$$0 \leq W_t \leq S_t$$

$$\sum_{i=1}^{n} q_{it} \leq \max\{(1-\delta_b)W_t - \delta_a, 0\}$$

With π_{it}, I_{t+1}, e_{it} defined as above.

Without in-stream demands the extraction problem is trivial (all flow is extracted). Note we are assuming spills are unavailable for extraction.

3.5.3 The parameterisation

To maintain generality, parameter ranges are specified rather than point estimates. Complete detail on the parameterisation is provided in appendix A.

Supply side parameters are based on statistics for 22 storages in the MDB. A data set on these storages was compiled from various sources including NWC (2011*a*), 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 primarily on an econometric analysis of irrigation farms in the southern MDB, drawing on ABARES survey of irrigation farms (Ashton and Oliver 2012).

Storage capacity K is the numéraire in parameterisation and is fixed at 1000 GL. Key parameter assumptions are summarised in table 3.1, for more detail see appendix A.

Table 3.1: Selected parameter ranges

	Min	Central case	Max
$E[I_t]/K$	0.23	0.71	1.18
c_v	0.40	0.70	1.00
ρ_I	0.20	0.25	0.30
$\alpha K^{2/3}/K$	0.03	0.09	0.15
δ_0	0.43	0.62	0.81
δ_{1a}	0.00	0.07	0.15
δ_{1b}	0.15	0.22	0.30
n	100.00	100.00	100.00
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

3.5.4 Solving the model

Use allocation

Given W_t and e_{it} the use allocation problem is

$$\max_{q_{it}} \sum_{i=1}^{n} \pi_{it}(q_{it}, I_t, e_{it})$$

$$\sum_{i=1}^{n} q_{it} \leq \max\{(1-\delta_b)W_t - \delta_a, 0\}$$

The first order conditions are then

$$\begin{aligned} P_t &= d_i^{-1}(q_{it}^*, I_t, e_{it}) \\ &= \max\{e_{it}(\theta_{h1} + 2\theta_{h2} q_{it}^* \mathscr{A}_i^{-1} + \theta_{h5} I_t),\ 0\} \end{aligned}$$

for all i, where P_t is the shadow water price.

Storage decision

With the use allocation problem solved by the above conditions, we can define the social welfare (planner's payoff) function Π

$$\Pi(W_t, I_t) = \sum_{i=1}^{n} \pi_i(q_{it}^*(W_t, I_t), I_t, e_{it}) \approx \sum_{i=1}^{n} \pi_i(q_{it}^*(W_t, I_t), I_t, 1)$$

where with large n the aggregate effect of the user productivity shocks becomes trivial and can be ignored. Figure 3.4 shows the implied planner's payoff function under the central case parameters. The steep portion is high reliability demand.

Figure 3.4: Social welfare, when $I_t = E[I_t]$ and $e_{it} = E[e_{it}] = 1$.

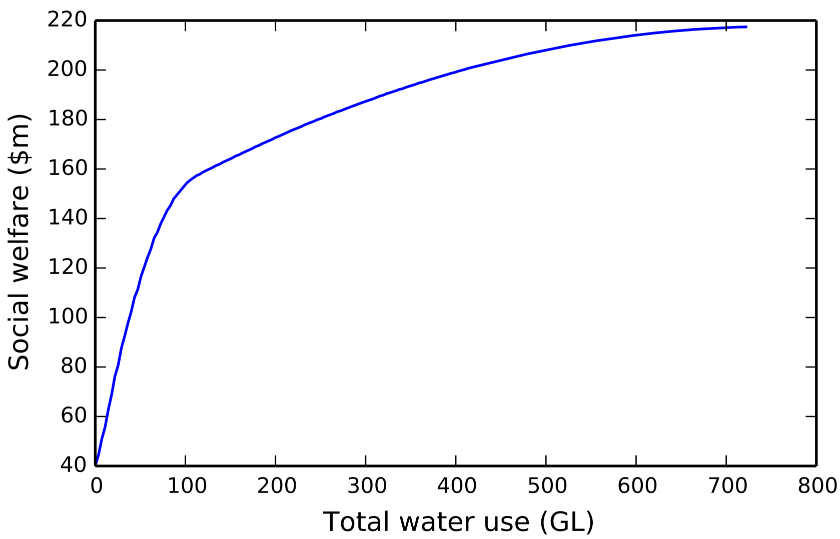

The storage decision is then

$$\max_{\{W_t\}_{t=0}^{\infty}} E\left\{\sum_{t=0}^{\infty} \beta^t \Pi(W_t, I_t)\right\}$$

subject to the storage transition rule, which is a dynamic programming problem with one policy variable W_t and two state variables S_t and I_t. We solve the problem numerically using policy iteration with continuous state and policy space. A form of tile coding (see section 8.4.2) is used to approximate the value and policy functions.

3.5.5 Results

Here, we compare optimal storage policy against a myopic release rule: $W_t = S_t$. Notionally, the difference between the optimal and myopic scenarios represents the 'gains from storage', or at least the gains from annual storage reserves. However, both of these scenarios remain somewhat stylised. Their main purpose is to provide benchmarks against which more realistic scenarios — with 'second best' property rights and / or imperfect policy rules — can be compared.

The myopic scenario might be interpreted as approximating a sub-optimal planner. Historically, storage management in the MDB has tended to be myopic. During the 1990s water was routinely 'over allocated': more water was allocated than could be used (on an annual basis) (MDBMC 1995). In practice, some unused water allocations remain in storage, so a myopic policy can still involve some annual reserves (see Brennan 2008*a*). For this reason, the 'no storage rights' scenario (NS) presented in chapter 5 is a more realistic depiction of myopic policy.

A more serious attempt to represent the behaviour of imperfect planners remains outside the scope of this thesis. Previous studies have found that planner storage rules tend to be myopic relative to optimal policy (Iglesias et al. 2007, Brennan 2008*a*, Hughes and Goesch 2009*b*). In appendix D, we consider scenarios where a planner adopts an SOP type storage rule[7].

Central case

Figure 3.5 compares the optimal policy function with a myopic policy for the central case parameters. Figure 3.6 shows the value function[8].

The optimal and myopic policies are then simulated for 500,000 periods. Figure 3.7 shows some sample results. The optimal policy accumulates storage reserves during wet periods in order to increase supply during dry periods, significantly decreasing the cost of drought.

[7] In these model runs we solve for the optimal $\bar{Q}$. This optimal SOP rule involves relatively negligible welfare loss compared with optimal policy. In practice, rules are unlikely to be optimally implemented given the information, incentive and policy uncertainty problems discussed above.

[8] The kink in the policy and value functions at low storage levels is caused by the fixed delivery loss

Figure 3.5: The policy function

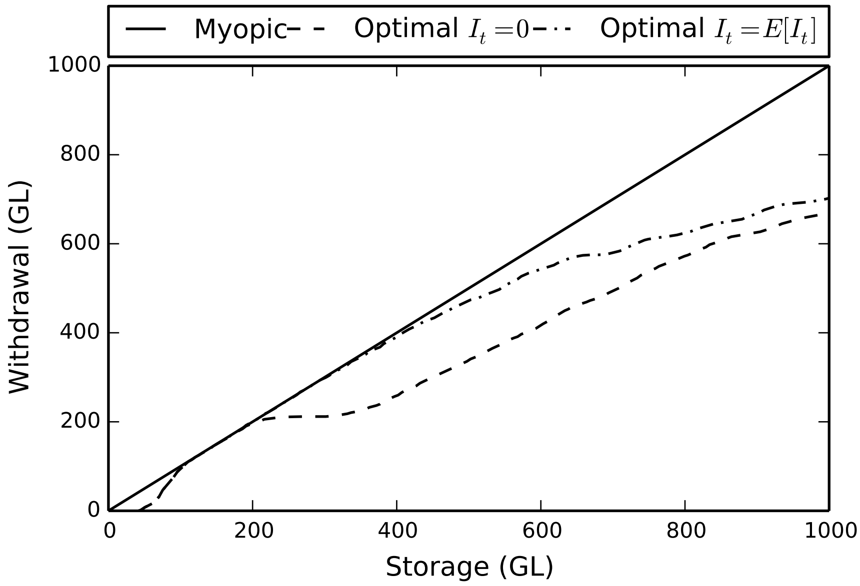

Figure 3.6: The value function

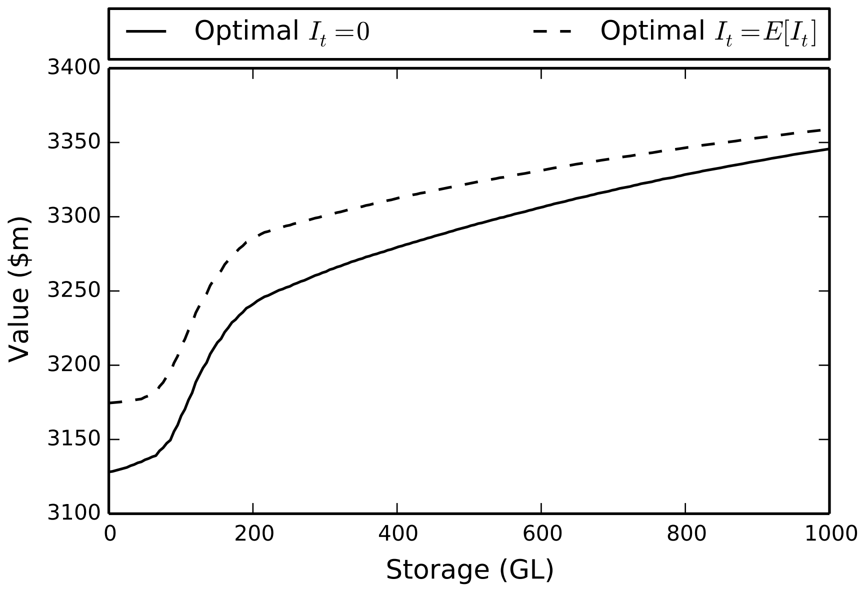

Figure 3.7: Sample simulation time series

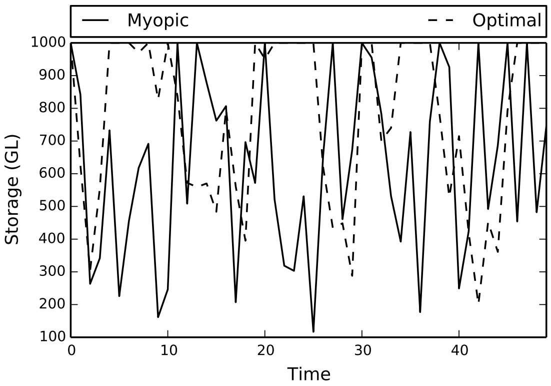

(a) Storage, S_t

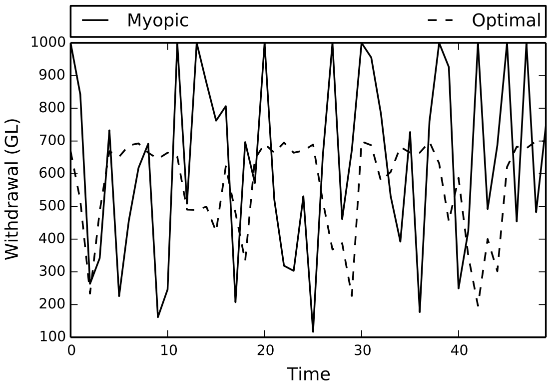

(b) Withdrawal, W_t

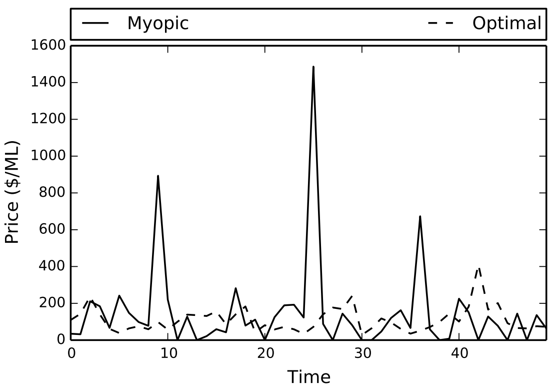

(c) Shadow price, P_t

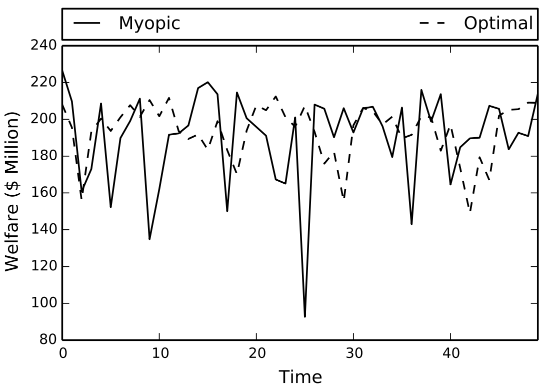

(d) Welfare, $\sum_{\mathcal{U}} u_{it}$

The central case results (table 3.2) are similar to existing studies: optimal storage provides a slight increase (around 3 per cent) in mean welfare. The welfare gains are all in dry periods (and the losses in wet years) so that a larger (around 25 per cent) decrease in the variance of welfare is observed.

Table 3.2: Central case results

	Mean	SD	2.5th	25th	75th	97.5th
Storage (GL)						
Optimal	696.92	282.88	159.76	459.98	1,000.00	1,000.00
Myopic	578.11	300.59	103.76	317.95	881.89	1,000.00
Social Welfare ($m)						
Optimal	186.72	25.25	128.73	178.74	202.29	209.84
Myopic	181.47	34.12	75.88	168.89	203.83	213.19
Withdrawals, (GL)						
Optimal	521.57	176.87	159.76	387.22	664.29	698.49
Myopic	578.11	300.59	103.76	317.95	881.89	1,000.00
Spills (GL)						
Optimal	130.29	279.77	0.00	0.00	87.13	1,054.54
Myopic	79.38	212.87	0.00	0.00	0.00	944.03
Shadow Price ($/ML)						
Optimal	165.79	278.91	28.75	64.88	168.71	878.66
Myopic	222.52	419.29	0.00	21.17	190.18	1,598.47

It is important to put these welfare effects in context. Firstly, these 'gains from storage' are comparable (if not larger) than the 'gains from trade' estimated using the same model (see chapter 6). The mean welfare effects in this type of model are always smaller than those of static models which typically present results only for dry scenarios. In wet periods where storages are overflowing and water has zero (or negative) value the gains from storage (and trade) are non-existent.

Further, the model may understate the cost of extreme drought years — given we do not account for the destruction of perennial crops or the loss of water for essential human needs — and extreme wet years — given we ignore flood damage.

A final point is that the low reliability (i.e. broadacre) profit functions include some dryland farm output, which dilutes the percentage welfare effects. For example, under the central parameters, the policy $W_t = 0$ still generates a mean welfare of $33m.

General case

Below, we solve the model for 1000 randomly drawn parameter sets. For each parameter set, we obtain the sample mean (table 3.3) and standard deviation (table 3.4) of key variables: S_t, W_t, $\sum_{\mathcal{U}} u_{it}$ and Z_t. For each parameter set and each statistic we also calculate indexes (myopic over optimal).

The results are summarised in figures 3.8 and 3.9. The average difference in mean welfare over the 1000 runs is 4 per cent, varying between 0 and 22 per cent.

Figure 3.8: Sample mean index (myopic over optimal)

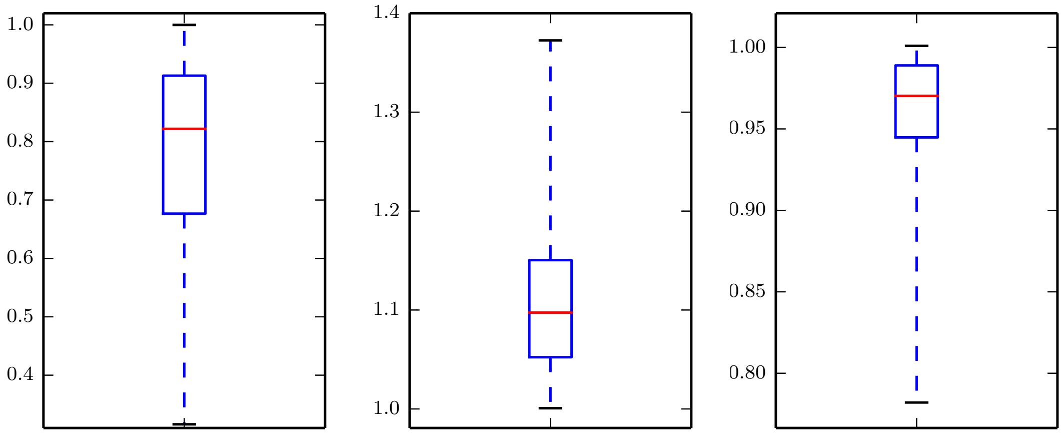

(a) Storage (b) Withdrawal (c) Welfare

Figure 3.9: Sample standard deviation index (myopic over optimal)

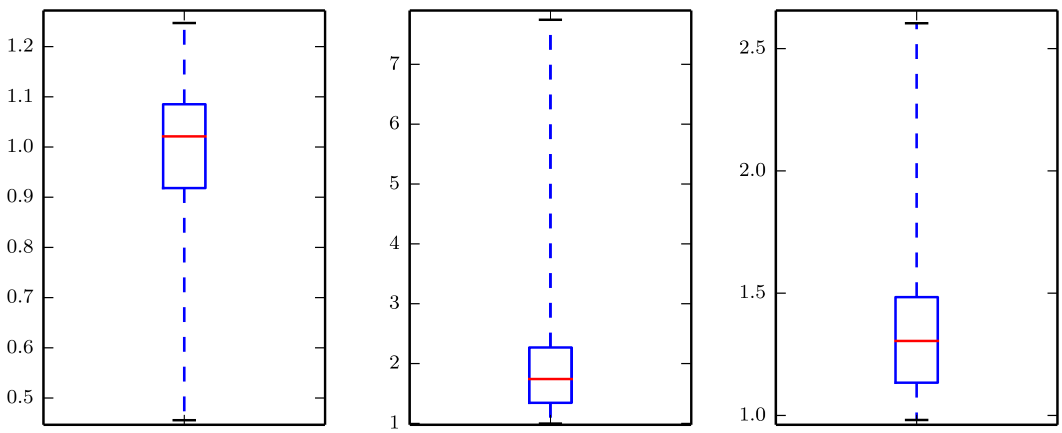

(a) Storage (b) Withdrawal (c) Welfare

Table 3.3: General case results — sample means

	Mean	Min	Q1	Q3	Max
Withdrawals (GL)					
Optimal	501.81	169.00	360.64	633.56	863.43
Myopic	547.90	213.32	421.23	682.10	866.72
Myopic / Optimal	1.11	1.00	1.05	1.15	1.37
Storage (GL)					
Optimal	682.19	351.32	628.56	750.91	913.87
Myopic	547.90	213.32	421.23	682.10	866.72
Myopic / Optimal	0.79	0.32	0.68	0.91	1.00
Social Welfare ($m)					
Optimal	186.19	51.07	125.21	235.59	421.04
Myopic	180.56	46.16	118.77	233.61	420.70
Myopic / Optimal	0.96	0.78	0.94	0.99	1.00
Spills (GL)					
Optimal	132.26	0.43	54.82	202.11	334.84
Myopic	92.47	0.00	17.40	160.31	256.79
Myopic / Optimal	0.56	0.00	0.31	0.78	1.00

Table 3.4: General case results — sample standard deviations

	Mean	Min	Q1	Q3	Max
Withdrawals (GL)					
Optimal	163.63	17.98	109.29	211.62	334.61
Myopic	270.34	106.76	236.43	308.62	348.26
Myopic / Optimal	1.93	1.00	1.34	2.27	7.74
Storage (GL)					
Optimal	273.45	154.24	255.11	295.26	335.38
Myopic	270.34	106.76	236.43	308.62	348.26
Myopic / Optimal	0.99	0.46	0.92	1.09	1.25
Social Welfare ($m)					
Optimal	26.80	2.32	14.84	36.70	95.06
Myopic	35.24	2.53	19.55	46.52	133.80
Myopic / Optimal	1.34	0.98	1.13	1.48	2.60
Spills (GL)					
Optimal	253.08	8.45	165.39	343.62	461.72
Myopic	195.65	0.25	91.21	298.16	384.29
Myopic / Optimal	0.70	0.00	0.55	0.86	1.00

Regression analysis

Next, we examine how the parameters are affecting the results. Here, we regress our mean welfare index against, the parameters, using the non-parametric 'random forest' method (see appendix B). Table 3.5 shows the 'importances': the contribution of each parameter to the explainable variation in the index. Figure 3.10 plots the effect of the key parameters, holding others constant at their central case values.

Table 3.5: Welfare index regression: importances and sample means

	Importance	$\leq$ 10th percentile	Full sample	$\geq$ 90th percentile
Index		0.88	0.96	1.00
c_v	33.93	0.89	0.71	0.55
$E[I]/K$	27.95	0.43	0.71	1.05
$\mathscr{A}_{low}$	14.92	1,466.30	3,972.57	7,669.18
N_{high}	9.36	57.22	50.32	45.86
δ_{1a}	8.53	47.5	53.0	66.7
α	1.44	9.70	8.97	8.73
ρ_I	1.32	0.25	0.25	0.25
δ_{1b}	1.29	0.22	0.22	0.24
δ_0	1.28	0.62	0.62	0.62

Figure 3.10: Welfare index regression results

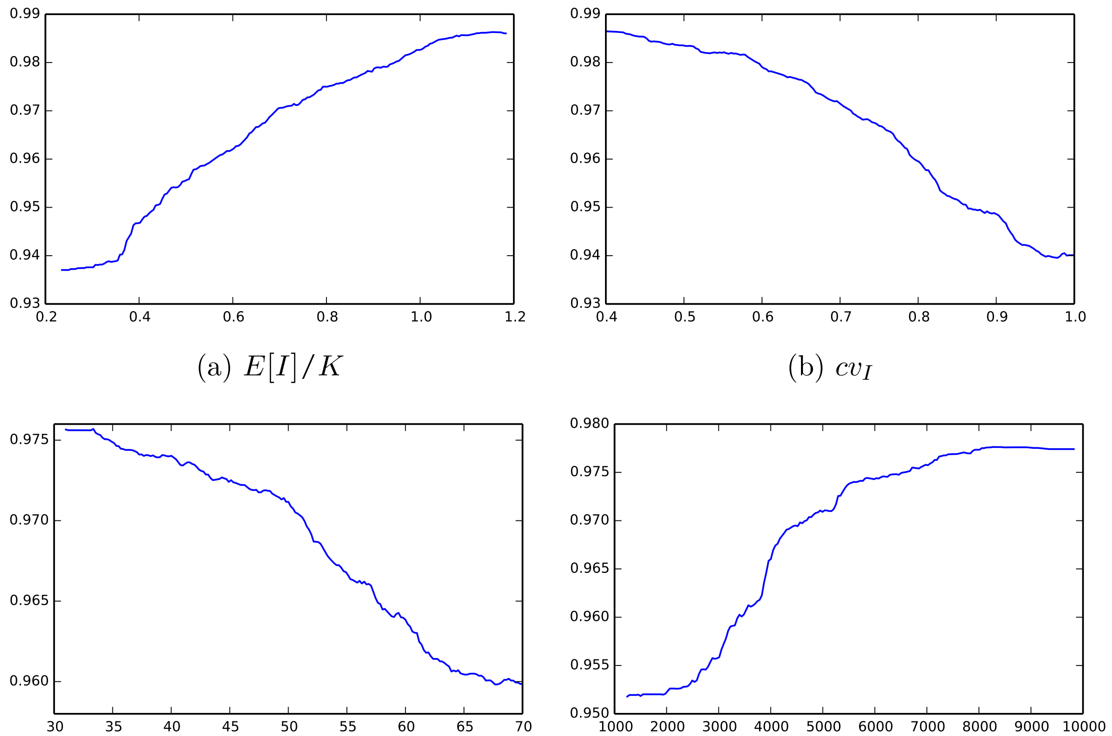

(a) $E[I]/K$ (b) cv_I

(c) n_{high} (d) $\mathscr{A}_{low}$

The two most important parameters are the ratio of mean inflow to storage capacity and the variance of inflow. The welfare gains from storage are much higher in rivers with high inflow variance and low mean inflow relative to capacity (since less water is lost to spills). A number of other expected results are observed, larger welfare gains are associated with: higher total water demand ($\mathscr{A}_{low}$), a higher proportion of high reliability demand (n_{high}) and higher fixed delivery losses (δ_a).

3.6 Conclusion

The water allocation problem in a regulated river can be divided into three parts: the storage decision (the inter-temporal allocation), the extraction decision (the trade-off between consumptive and in-stream use) and the use allocation (the spatial allocation).

Economists have focused mostly on use allocation. Here, economists generally argue for markets over government control. Until recently, storage and extraction decisions have been the domain of engineers. While inter-temporal issues have received attention from economists in other natural resources they have often (with a few notable exceptions) been neglected with surface water.

In this chapter, we presented a parametric version of a planner's water allocation problem in a regulated river. Our numerical results largely confirm previous studies, with optimal storage policy leading to a gain in mean welfare (on average 4 per cent) and a larger reduction in the variance of welfare, compared with myopic policy. We also detail how these effects vary with the parameter values. In rivers with low and highly variable inflows, storage policy has a larger effect on welfare.

Economists have long documented the failures of government management of natural resources, water in particular. Here we have considered three main limitations of government. First, governments have incomplete information on the preferences of water users. Second, there is a risk of political bias. Third, governments can create policy uncertainty, by failing to commit to or communicate complete water allocation rules.

The discussion of government in this chapter has been brief. While economists often focus on the failures of government, it is important to acknowledge that bureaucratic structures (e.g., hierarchies, independent agencies and policy rules) evolve in response to these problems. We consider some of these institutional issues further in the next chapter.

Chapter 4

Water property rights and markets: theory and practice

4.1 Introduction

Water property rights have a long history in Australia and the western US. US water rights can be traced back to the prior appropriation system, which evolved during the 1800s. Australian water rights can be traced back to the Victorian Irrigation Act of 1886. In both regions, formal water rights predate the development of organised market trade by some time.

In Australia, early water market reforms began during the 1980s (NWC 2004). Today, the Southern MDB is viewed as perhaps the most successful water market in the world (Grafton et al. 2011*b*). While market reforms have been less comprehensive in the US, trading is extensive in a number of regions including California (Hanak and Stryjewski 2012).

Market reforms do not preclude government involvement in water allocation. Firstly, much government effort is involved in defining and enforcing water rights and in regulating water trading. Beyond this however, governments typically maintain significant control over water allocation within regulated rivers, particularly over storage and extraction.

Institutional arrangements around water allocation are notoriously complex, often involving a seemingly chaotic mix of government control, decentralised markets and collective arrangements. This complexity has elicited much interest from the various forms of Institutional Economics (IE). In particular, water has been central to the research of Ostrom (2010*b*) and the concept of 'polycentric' governance.

This chapter begins by defining water property rights and markets in the context of our water allocation problem from chapter 3. Here, we introduce a standard approach to water property rights in regulated rivers, common in Australia and the US, and the concept of a water spot market. We then provide an introduction to the literature on water markets.

Our focus then turns toward reconciling our abstract treatment of water property rights with the more complex reality. To begin, we consider some of the theory on property right institutions for natural resources. Then we discus how complex 'polycentric' water institutions can be viewed as hierarchies of property rights. Finally, we describe the water right and market systems of the MDB and the western US.

4.2 A standard water right and market system

Below we outline a standard water property rights system broadly reflective of that adopted in the major regulated rivers of both the MDB and the western US[1]. Under this 'release sharing' (Dudley 1988*a*) approach, storage and extraction policies are centrally determined and users receive shares in releases. For now we assume a standard storage rule

$$W_t = \min\{S_t, \bar{S}\}$$

$$\bar{S} \leq K$$

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

The storage and extraction polices define an aggregate 'allocation' A_t and a maximum allocation $\bar{A}$

$$A_t = E_t - \mathcal{L}_E(E_t)$$

$$\bar{A} = \bar{S} - \mathcal{L}_1(\bar{S}) - \mathcal{L}_E(\bar{S} - \mathcal{L}_1(\bar{S}))$$

$$0 \leq \lambda_i \leq 1, \sum_{i=0}^{n} \lambda_i = 1$$

[1]For more detail on the MDB and western US, see section 4.5. This approach was standard in most MDB regions until the introduction of storage rights, see chapter 5.

Typically, use allocation involves priority classes. With two priority classes (as in section 3.4.3) the water allocation a_{it} of user i in period t is given by

$$a_{it} = \begin{cases} \min\{\lambda_i A_t, \lambda_i \bar{A}\ \} & \text{if } i \in \mathcal{U}_a \\ \lambda_i(\max\{A_t - \sum_{\mathcal{U}_a} \lambda_i \bar{A}, 0\}) & \text{if } i \in \mathcal{U}_b \end{cases}$$

At this point, water property rights and water allocation rules appear equivalent. However, there are two key distinctions. Firstly, property rights convey a degree of certainty, that is the shares λ_i are fixed: or at least can not be altered without compensation. Secondly, property rights offer some autonomy to users, including the ability to trade water.

4.2.1 The water spot market

A water spot market involves user payoff functions

$$u_{it} = \pi_{it}(q_{it}, C_t) + P_t(a_{it} - q_{it})$$

where P_t is the market price for water. The users' problem is to choose q_{it} to maximise profit given a_{it}, C_t and P_t

$$\max_{q_{it}} u_{it}$$

With perfectly competitive markets the equilibrium price satisfies

$$d^{-1}(q_{it}, C_t) = P_t \quad \forall i$$
$$\sum_{i=1}^{n} q_{it} = \sum_{i=1}^{n} a_{it}$$

Here the market achieves optimal use allocation. Where storage and extraction policies are also optimal we replicate the social planner's benchmark.

Transfer costs

In practice, trade in water is not costless. A positive transfer cost can be represented as a tax τ on trade (in this case on buyers):

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

Defining $p_{it} = d_i^{-1}(a_{it}, C_t)$ as user i's marginal value (willingness to pay) for water before trade, the clearing price now satisfies

$$q_{it} = \begin{cases} d_i(P_t, C_t) & \text{if } p_{it} \leq P_{it} \\ a_{it} & \text{if } P_t < p_{it} < P_{it} + \tau \\ d_i(P_t + \tau, C_t) & \text{if } p_{it} \geq P_{it} + \tau \end{cases}$$

4.3 Water market literature

The economic literature on water rights and markets is vast. We provide only a brief introduction here. For more detailed reviews see Dudley (1992) or Chong and Sunding (2006).

Early studies highlighted government inefficiencies and made the case for markets. Hirshleifer et al. (1969) documented large differences in water prices between regions and user groups in the US. Burness and Quirk (1979) detailed problems with individual priority ordering (under US prior appropriation). Randall (1981) highlighted inefficiencies in the Australian case and made similar arguments for trade.

All of these studies were careful to acknowledge the limitations of water markets, particularly the problem of externalities. Water trading across space can have a range of external effects (aka third party effects) on other water users due to: return flows, delivery losses and constraints, in-stream values, water quality effects among other causes.

Much of the recent literature focuses on these spatial externality problems (see for example Brennan and Scoccimarro 1999, Heaney et al. 2006). The central question is whether these effects are significant enough to warrant restrictions on trade or changes to water property rights. Heaney et al. (2006) argue that many externalities are too minor to warrant intervention, while Brennan and Scoccimarro (1999) stress the need for empirical estimates.

The US literature on water trade externalities focuses largely on return flows, see for example (Howe et al. 1982; 1986). A commonly proposed reform (see Johnson et al. 1981) is to define property rights to water net of return flows (assuming they are measurable)[2].

There is also much empirical work on the evolution of water markets. Anderson (1961) documents one of the earliest examples of a water market: the US South Platte basin. The region has since been studied by many others (Carey and Sunding 2001). Hanak and Stryjewski (2012) provide much detail on Californian markets while Libecap (2011*b*) provides a broader US perspective. Australian water markets have been well documented by the NWC (2011*a*) and others (such as Brennan 2006). Grafton et al. (2010) provides a comparison of Australian and US water markets.

We consider other strands of the literature in later chapters, specifically: storage issues in chapter 5; priority rules in chapter 6 and in-stream demands in chapter 7.

4.4 Property right institutions

The economic literature on water markets frequently adopts the conventions of the 'property rights school'[3]. Here property rights exist to 'internalise externalities' and emphasis is placed on the efficiency of markets in contrast with government control or open access. These views were recently summarised by Libecap (2009):

> In one way or another, all environmental and natural resource problems associated with overexploitation or under provision of public goods, arise from incompletely defined and enforced property rights. As a result private decision makers do not consider or internalize social benefits and costs in their production or investment actions. (Libecap 2009; pp. 129)

Many object to this property rights view of the world, arguing it can not account for the institutional complexity observed in natural resources. Water has long been a primary example, as Vincent Ostrom observed:

> Few areas of American political and economic experience offer a richer variety of organizational patterns and institutional arrangements than

[2]As Griffin and Hsu (1993) note this approach is based on an oversimplification of the problem. In practice, given the spatial complexity of return flows a simple net water use right will not lead to a first best outcome.

[3]Baland and Platteau (1996) use the term 'property rights school' in reference to a body of literature (for example Hardin 1968, Demsetz 1967) promoting private property approaches to natural resource allocation.

> the water resource arena. Yet these patterns of organization have developed in a way that seems to conform neither to the prescriptions of political scientists nor to the prescriptions of economists. (Ostrom 1962; pp. 450)

Much of the early debate centred on the distinction between open access and common property.

An early point of contention, was the assumption that an absence of government control or government enforced rights necessarily implied an open access outcome[4]. Here water was central in the development of common property concepts, particularly in the research of Elinor Ostrom[5]. There now exists vast theoretical and empirical literature on common property in natural resources (Baland and Platteau 1996).

However, property rights theory is still considered insufficient to explain many aspects of institutional systems. In response, there is a growing body of institutional economics (IE) literature, including much specific to water (for example Ostrom et al. 1994, Saleth and Dinar 2004, Challen 2000, Sharma 2012). Methodological approaches vary considerably, often landing somewhere in-between New Institutional Economics (NIE) — with its emphasis on neoclassical theory — and old IE — with its more anti-theoretical descriptive approach (Saleth and Dinar 2004).

In this section, we attempt to reconcile these broader views with the reductive approach adopted in this thesis. We begin by defining some key concepts: transaction costs, property rights, externalities and common property.

4.4.1 Transaction costs

Within NIE the term transaction cost has a very broad meaning. For example Cheung (1998) defines transaction costs as: "...just about all the conceivable costs in society except those associated with the physical processes of production or transport." (Cheung 1998; pp. 515)

In this thesis we draw a distinction between transfer costs and institution costs. Transfer costs are associated with negotiating, processing and enforcing individual transactions (e.g. water trades). Institution costs refer to transaction costs that

[4]This assumption was made most famously by Hardin (1968) who popularised the 'tragedy of the commons'. Subsequent to this, the terms open access and common property were often confused in the literature, see Quiggin (1993).

[5]Elinor Ostrom's thesis documented institutions for groundwater management in California (Ostrom and Ostrom 1972).

are not marginal to particular transactions (fixed transaction costs). Institution costs include the costs of establishing and enforcing property rights and in our context the costs of water accounting and measurement systems.

4.4.2 Property rights

The economic concept of property is extremely general, referring to the practical ability of any entity to control or use a good. Economic property rights are distinct from legal property rights in extending to rights enforced by non-government organisations, even rights enforced by informal rules or social norms (Eggertsson 1990).

Property rights are commonly divided into three components or 'powers' (Scott 2001): the right to to use a good, the right to earn income from the good and the right to transfer ownership of the good to others. With natural resources it is common for property rights to exist without transfer rights (Ostrom 2010*a*).

Property rights are commonly viewed as a bundle, both in the sense of conveying multiple powers and in the sense of containing many distinct rules for specific components or uses for a good. In the case of water, distinct rights may be defined to storage, delivery, extraction and use.

Scott (2001) describes property rights for natural resources in terms of six characteristics: exclusivity, duration, flexibility, quality of title, transferability, and divisibility. Scott (2001) imagines property rights varying continuously along each of these dimensions, with rights said to be more complete (less attenuated) when high scores are attained on each. In this thesis, the focus is primarily on the exclusivity characteristic.

4.4.3 Externalities

By externality, we refer to any effect of an agent on the welfare of another agent, external to the price mechanism (the 'Pareto-relevant' externalaties of Buchanan 1962).

While this notion of externality is conventional it is not without controversy (see for example Dahlman 1979, Randall 1983; 1974). Part of the problem is that externalities are always a symptom of some other fundamental cause. Following Randall (1983) we consider two types of externality — non-exclusiveness and non-rivalry — with two fundamental causes — transaction costs and information problems.

The classic example of non-exclusiveness is the open access resource: in the absence of property rights, users ignore how their extraction affects others resulting in over exploitation: 'The Tragedy of the Commons'.

In theory, this type of externality can be eliminated with the definition of exclusive rights. In practice, non-exclusiveness persists because of transaction costs, specifically the costs of defining and enforcing exclusive property rights. The unique physical characteristics of fugitive natural resources — such as water — mean that even where property rights are well defined they are typically only ever partially exclusive (Eggertsson 1990).

Like exclusiveness, rivalry is best viewed as a continuum. Degrees of non-rivalry are frequently encountered with natural resources. Water examples include non-consumptive flows and storage and delivery capacity (which are non-rival below capacity constraints). Generally externalities caused by non-rivalry can not be eliminated from markets, regardless of the degree of exclusiveness, because of information problems (see for example Oakland 1974).

4.4.4 Common property

Common property refers to the idea of community ownership of a resource: where "a community controls access to a resource, by excluding outsiders and regulating its use by insiders" (Eggertsson 1990; pp. 36). Literature on common property involves both empirical studies (see Ostrom et al. 1994) and theoretical models (e.g., cooperation as an equilibrium outcome in repeated games, see Baland and Platteau 1996).

Common property generally implies the existence of local institutions: that is 'internal access rules' designed and enforced by resource users. In most cases, these rules are easily interpreted as local property rights. This is best understood by example.

Ostrom et al. (1994) document the results of 47 case studies of irrigation schemes, including 29 farmer managed and 14 government managed systems. All but three of the farmer managed schemes adopt explicit water allocation rules (these exceptions are all rated as low performing).

Three types of rules are observed: 'fixed percentage shares' (similar to section 4.2), 'fixed orders' and most frequently 'fixed time slots', where farmers must extract water during designated periods. As Ostrom et al. (1994) note, time slot rules can easily be understood in terms of monitoring and enforcement costs.

There is much debate over when local institutions might be preferred to government. Attention is usually placed on the number of resource users, their homogeneity and the availability of information (Ostrom et al. 1994, Baland and Platteau 1996). A popular, if unsatisfying conclusion, is that some mix of government and local institutions will generally be ideal (Baland and Platteau 1996).

4.4.5 The evolution of property rights

A theory of the emergence of property rights begins with Demsetz (1967):

> "...property rights develop to internalize externalities when the gains of internalization become larger than the cost of internalization". (Demsetz 1967; pp. 350)

According to this theory property rights always evolve in response to needs: for example where changes in preferences or technology increase the scarcity value of a resource, or decrease the costs of exclusion. In the case of water there are endless examples of property rights emerging in response to increases in scarcity. Some specific Australian and US examples are discussed in section 4.5.

Scott (2001) elaborates more on the mechanisms involved, suggesting a process of interaction between property rights suppliers and demanders. We generally expect governments to be the suppliers, but non-government organisations can also have prominent roles (as is the case with common property).

The idea that property rights reforms are driven by end users is consistent with observation. With US prior appropriation rights, governments formalised a system established independently by users (Scott 2001). While water markets in the MDB flourished after state government agreements in the 1990s, there are reports of unofficial trading as early as the 1940s (NWC 2011*c*).

The theory of Demsetz (1967) is sometimes viewed as 'naive' (Eggertsson 1990), in not accounting for the incentives of property rights suppliers. The alternative 'interest-group theory' emphasises the role of rent seeking. It suggests welfare enhancing property rights reforms will not occur, where adversely effected parties can effectively lobby governments (Libecap 1978).

The interest group theory has two main implications. Firstly, prevailing property rights may not be ideal. Secondly, property rights reforms will require compensation packages to be politically feasible.

4.4.6 Water institutions as property rights hierarchies

While there can be a reluctance to impose much structure on these complex 'polycentric' systems, it is acknowledged they often take the form of hierarchies (North 1990). The pronounced hierarchical structure of water institutions has been recognised by many economists (Ciriacy-Wantrup 1967, Ostrom et al. 1994, Challen 2000, Saleth and Dinar 2004).

Challen (2000) describes the water institutions of the MDB as hierarchies, where each level corresponds to one of: state property, private property, common property or open access. Here, we generalise this idea, and describe water institutions simply as hierarchies of property rights.

The idea is best understood by example (section 4.5 provides further detail for Australia and the US). Government is typically at the top of the hierarchy. The highest level of water rights are those of either nations or states, often over cross border resources (e.g., the rights of NSW, VIC and SA to the Murray River).

Jurisdictions are then able to implement there own subsidiary water rights systems. In the US the rights defined by state governments are often held by collectives (e.g. irrigation districts) or by other government agencies, who in turn can define further rights systems, and so on.

A hierarchy can emerge where rights holders (be they governments, collectives or individuals) have the ability to define subsidiary property rights. These subsidiary rights are more than a subdivision of existing rights. Rather, rights holders take responsibility for defining and enforcing new property rights systems (as in the common property examples above).

The characteristics of these new rights can differ significantly from their parent rights. Rights at the higher levels are often more completely (i.e. exclusively) defined. Higher level rights also tend to be more secure than lower rights. However, frequently the nature of the rights held by states or nations are similar to those of end users (see section 4.5).

The obvious questions are why and how these hierarchies evolve. Some potential explanations can be drawn from the disparate economic literature on hierarchies[6]. The most obvious is information. Property right hierarchies allow both resource allocation decisions and the provision of institutions to be devolved to entities with more local information.

[6]Discussion of hierarchies is mostly found in the literature on federalism in government (see for example Oates 1999) and in the theory of the firm (Williamson 1973, Stiglitz 1975, Radner 1992, Grant 1996). A key exception is the paper of Simon (1965).

These complex hierarchies can also be viewed as a product of evolutionary processes. For example Simon (1965) argues that complex systems tend to evolve through hierarchical combination of self sustaining local systems. This suggests a degree of path dependence in water institutions: where local institutions arise to meet the needs of the time, but then become difficult to remove. Both Hanemann (2014) and Libecap (2011*a*) argue that the persistence of local water institutions in the US (e.g., irrigation districts) is largely the result of path dependencies.

The idea of property rights hierarchies may seem a touch semantic, but it serves a number of purposes. Firstly, it shows that the design of water institutions is largely about the design of property rights. Whether we are interested in formal arrangements among nations or informal arrangements among farmers, similar economic concepts (i.e., externalities and transaction costs) can be applied. Secondly, it shows how government enforced property rights can coexist with informal local rights and with government control[7]. Finally, it provides a path towards modelling these systems.

4.5 Water property rights in the MDB and western US

There is vast economic, legal and historical literature documenting the evolution of water property rights both in Australia (Davis 1967, Challen 2000, Crase 2008) and in the United States (Hanemann 2014, Libecap 2011*b*, Donohew 2009, Griffin 2012, Getches 1984).

We do not attempt a comprehensive summary here. Rather, we focus on a few issues of interest to this thesis, such as: storage, water trading and rights hierarchies. While there is a long history of comparison between the western US and Australia on water, there are few studies comparing the regions property rights (exceptions include Davis 1967, Grafton et al. 2010).

4.5.1 Water property rights in the MDB

MDB water rights can be traced back to the Victorian irrigation Act 1886 and the NSW Water Act 1905, which followed Deakin's 1884 Royal Commission. This legislation replaced riparian water rights with statutory water permits (Davis 1967). These changes were motivated by plans for large scale irrigation development (i.e., large dams) and the related policies of closer and soldier settlement (Davis 1967).

[7] Dudley (1992) made a very similar point in the context of capacity sharing: groups of users could own a collective capacity share and then define their own internal institutions.

Very early state governments established rules for allocating water against these permits during shortages (Davis 1967): recognising that even with dams, supply would remain variable. In general, these rules amounted to proportional sharing of available water, subject to a small number of priority classes[8]: our standard property rights of section 4.2.

Under these systems, storage, extraction and use decisions were all made by government. Only in recent times have reforms enabled trading and decentralised storage (see chapter 5). State storage management policies have evolved in line with irrigation development: Victoria has more conservative storage policy — given its focus on dairy and perennial crops — while NSW is more liberal — consistent with its focus on annual crops (Hughes et al. 2013).

The property rights of the states (NSW, VIC and SA) to the Murray River, were defined by the River Murray Waters Agreement of 1915[9]. The bargain established in 1915 remains largely unchanged: NSW and Victoria receive 50 per cent of Murray River flow (at Albury) and ownership of their tributary flows, subject to providing a fixed annual flow to SA (Crase et al. 2004)[10]. The current version amounts to a capacity sharing system: NSW and Victoria have percentage shares of inflow and storage capacity in the two main storages (Hume and Dartmouth).

Water trade in the MDB

The MDB water market is frequently considered the most advanced in the world. By international standards the southern MDB spot market (i.e., annual allocation market), involves low transaction costs and high trade volumes (Grafton et al. 2011*b*). Much user level trading occurs both within and across irrigation areas and across state boundaries (NWC 2011*b*). Trade is facilitated by private brokers and private electronic exchanges.

Early steps towards markets began during the 1980s (NWC 2011*c*). Initial reforms followed a gradual user-driven process (NWC 2011*c*). The pace of reform then increased with state government agreements in 1994 — The Water Reform Framework — and 2004 — The National Water Initiative — (NWC 2004).

Trade volumes have increased dramatically since the 1980s, with the removal of restrictions on trade, reductions in transaction costs and increasing water scarcity

[8]In NSW and VIC priority has long been given to perennial horticulture over annual broadacre cropping (Davis 1967; pp. 674).

[9]Negotiations towards the agreement began as early as 1863. The current iteration of the agreement is known as the Murray-Darling Basin Agreement (MDBA 2014*d*).

[10]Under normal — non-drought — conditions the upper basin states (NSW and VIC) have to allow at least 1850 GL to reach SA. NSW and VIC then receive the 50 per cent each of the remainder of Murray inflow (above Albury).

Figure 4.1: Southern MDB water trade volumes

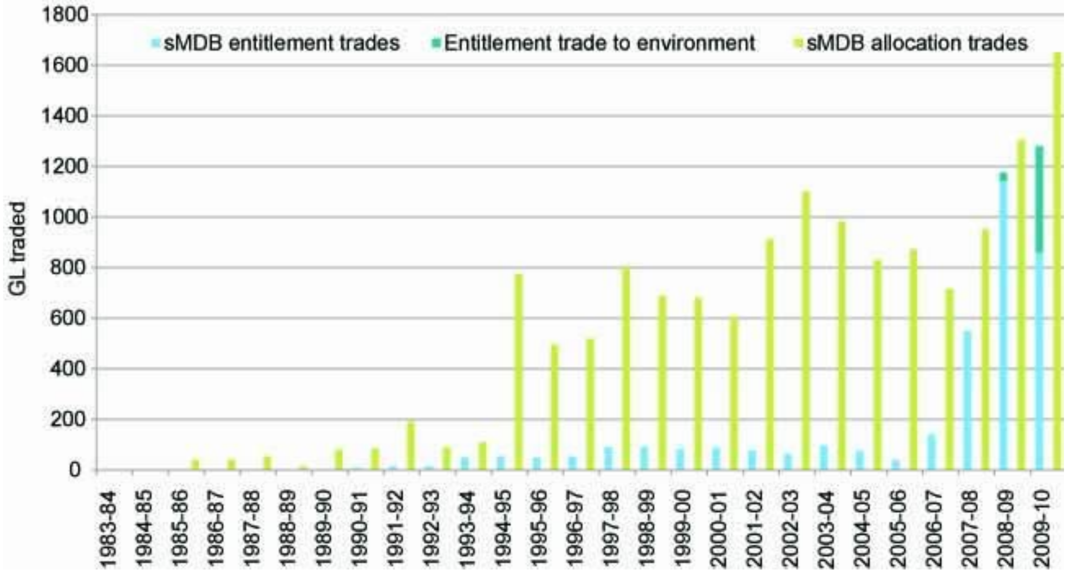

Source: NWC

Figure 4.2: Southern MDB water allocation prices

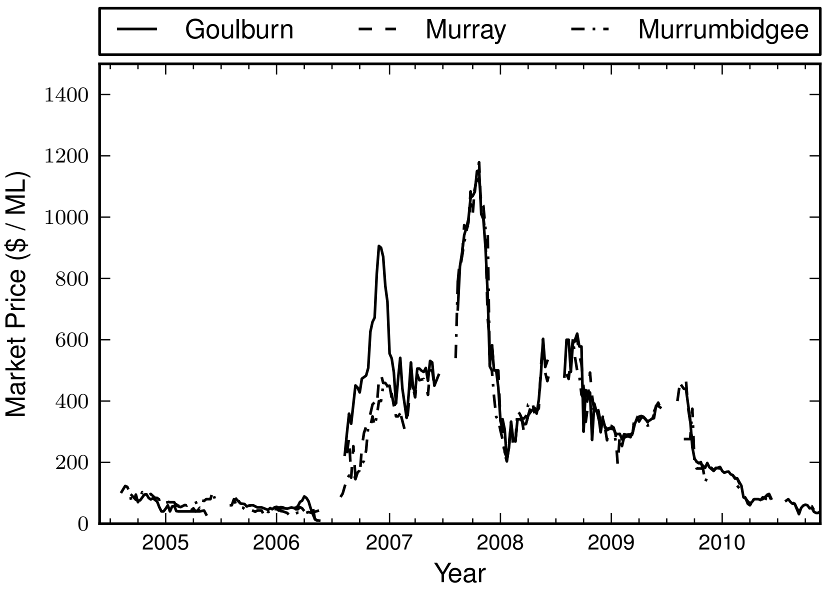

Source: waterexchange.com

(figure 4.1). Trade volume in the allocation market is much greater than in the permanent entitlement (water share) market, given lower transaction costs.

The southern MDB water market played a well documented role in reallocating water during the recent drought. The peak of the drought saw unprecedented increases in prices (figure 4.2) and the reallocation of large volumes of water across irrigation sectors — from broadacre to horticulture — and across state boundaries — from NSW into SA — (NWC 2011*b*).

Water trade also occurs in the northern MDB. Although trade volumes are smaller and markets are more localised given limited hydrological connectivity (NWC 2011*b*).

4.5.2 Water property rights in the western US

Prior appropriation

Discussion of US water rights usually centres on the doctrine of prior appropriation: a system of property rights which evolved in the western US during the late 1800s[11]. Prior appropriation involves two main principals: 'beneficial use' and 'first in time, first in right'. Firstly, to establish and maintain a right, water has to be consistently used for some 'beneficial' purpose (i.e., use it or lose it). Secondly, water allocation follows a priority ordering based on the dates rights were first established.

The principles of prior appropriation can be seen as adaptations to the prevailing conditions: rapid decentralised development of unregulated rivers. During this period, the main externality problem was the effect of new developments (i.e., extraction) on existing downstream users. Here, date ordering provided investment confidence (Burness and Quirk 1979): it ensured investors were protected against future upstream development.

Western US water rights and laws have evolved in a very gradual decentralised way. While much new law has been introduced since the 1800s, rarely has this completely displaced existing law[12]. As a result of this complex legal environment, US water rights are generally less secure than Australian rights and legal disputes are frequent.

[11] There are differing historical accounts of the evolution of prior appropriation. A popular view is that the system originated in California in the years following the 1859 gold rush (Tarlock 2002).

[12] For example, in California and Texas, aspects of the older riparian law remain in force (Libecap 2011*b*).

This legal complexity is partly a consequence of the historical development of the western US. Major water extraction began during the 1800s, before the construction of large dams (Dudley 1992) and even before the formation of state governments. The construction of dams then occurred in a decentralised way, with multiple federal and state agencies independently pursing their own projects.

While aspects of the prior appropriation doctrine remain entrenched in US water law, their practical influence is often overstated. For one, prior appropriation principles are often not strictly enforced (Tarlock 2002). For example, during droughts governments often ration water on a proportional basis, rather than strictly following the date order. Further, the majority of water users — those receiving water from 'projects' (i.e., large dams) — hold more modern water rights in the form of long-term contracts. As Tarlock (2002) explains:

> Dams made it increasingly unnecessary to enforce water rights in the rigorous manner that the doctrine suggests and helped produce the culture of non-enforcement of the beneficial use doctrine. The threat of priority enforcement decreased substantially. Water rights became more of a general water entitlement to use water rather than the right to a specific quantity used in a non-wasteful manner as specified by the formal doctrine. As a result, prior appropriation became more and more of a shadow doctrine. Increasingly, the federal government became the water master on large rivers such as the Colorado, Missouri, and Columbia. Federal and state contracts often became the real allocation rules (Tarlock 2002; pp. 771)

Water 'contracts'

US water property rights are best understood as a hierarchy. State government agencies — such as the State Water Resource Control Board (SWRCB) in California — define the formal water rights. Most of these rights are based on the prior appropriation system: at least in theory.

However, for most large dams only aggregate water rights have been established. These rights are held by a 'water master', typically a state or federal agency depending on who funded construction. In California, the federal Central Valley Project (CVP) is managed by the United States Bureau of Recalamation (BoR) and the California State Water Project (SWP) is managed by the Californian Department of Water Resources (DWR). These agencies have responsibility for storage management decisions, subject to strictly enforced flood mitigation rules.

Water masters define their own water rights, in the form of long-term contracts. These 'contracts' — typically held by irrigation districts — operate similar to Australian style water rights: allocations are made proportionally subject to a small number of priority classes and trade between contract holders is possible. Contracts have fixed terms of 20 to 50 years, but are often renewed with minimal change[13].

This hierarchical structure allows water rights to evolve without state government involvement. An instructive example is that of the Texas river authorities, such as the Lower Colorado River Authority (LCRA)[14]. For some time the LCRA has been purchasing old prior appropriation water rights, then offering new contracts against them. This process effectively retires these legacy water rights and replaces them with new rights more adapted to the current environment[15].

Irrigation districts

Annual water contracts are typically held by collectives, particularly irrigation districts (ID) that represent large numbers of irrigation farmers. The IDs then establish their own local rights systems and have responsibility for their monitoring and enforcement.

Approaches vary across districts but are often similar to our standard water rights system (section 4.2). For example, the Westland's ID in California, implements a proportional use rule, all be it with some additional priority classes. These allocations can be traded internally with few constraints.

The governance structure of IDs also varies. Westlands has an elected board, votes are weighted by member land holdings, in other districts voting is unweighted. Some IDs resemble jurisdictional governments with broad powers including the ability to tax all residents even non-farmers (Libecap 2011*a*).

Collective institutions have also evolved for urban water users, such as the Metropolitan Water District of Southern California (MWDoSC), that represents 26 urban water utilities, including Los Angles and San Diego (Zetland 2008).

[13]In California most CVP contracts have recently been renewed with near identical terms, although contract renewal is not always a formality (Tarlock 2002).

[14]LCRA is a state energy and water utility with responsibility for the Texas-Colorado River. Similar to BoR LCRA have aggregate water rights, associated with major storages, however there remain a large number of small prior appropriation rights established prior to construction of major storages.

[15]This practice of gradually acquiring water rights is seen by some economists as anti-competitive behaviour (Griffin and Characklis 2002). The more positive interpretation presented above was confirmed by the LCRA (2013, pers. comm., Ron Anderson, 9 May). If authorities like LCRA did more to encourage trading among contract holders this might allay competition concerns.

One disadvantage of this complex hierarchy is that trading can be difficult. While trade within IDs can occur relatively easily, trade between districts is complex as it requires group decision making[16] and is subject to water master approval. Trading across projects is more difficult as it is subject to state government approval. Interstate trade is virtually impossible. As such, most water trading is localised (Libecap 2011*b*).

Trade approvals are stricter and more costly in the US than Australia (Grafton et al. 2011*b*). Firstly, US regulators tend to put greater emphasis on preventing external effects from trade[17]. Secondly, constraints on trade arise from environmental legislation, particularly the Endangered Species Act[18]. Thirdly, in the US trades are approved on a case-by-case basis, rather than through ex ante trading rules, as in the MDB.

Hanak and Stryjewski (2012) argue, that despite these constraints, the Californian water market is still reasonably efficient. Figure 4.3 shows Californian water trade volumes. While there is much evidence to suggest Australian water markets are superior (Grafton et al. 2011*a*), there are a number of counter arguments.

Firstly, a general comparison between the US and the southern MDB maybe unfair. There are a number of examples of US regions, such as the South Platte Basin and the Texas Lower Rio-Grande, with very well defined, homogeneous, user level water rights, reminiscent of the southern MDB. In both cases, property right reforms — which date back to the 1950s — have fostered efficient markets.

Secondly, users have found ways to adapt to the high transaction costs. Most trades are large permanent transactions, negotiated between collective organisations (Hanak and Stryjewski 2012). Dry year options contracts are also common (Hanak and Stryjewski 2012). In drought periods, government 'water banks' have evolved, purchasing water in low value regions and selling into high (Hanak and Stryjewski 2012).

Thirdly, trade is open to urban users. In fact, much of the trade in California has been from agriculture to cities (Hanak and Stryjewski 2012). In Australia, trade

[16]For example, Westland runs a 'trading pool': members contribute funds to the pool, the Westland's board enters the market with these funds and then distributes its purchases in proportion to member contributions.

[17]The main issue here is return flows. In many heavily exploited US rivers, changes in return flows can affect downstream consumptive users. Generally in the MDB, changes in return flows only affect in-stream flows.

[18]A prominent example in California are limits on pumping from the Sacramento-San Joaquin River Delta, to protect endangered fish. This constraint severely limits trade of water from North to South (Hanak and Stryjewski 2012).

Figure 4.3: Californian water trade volumes

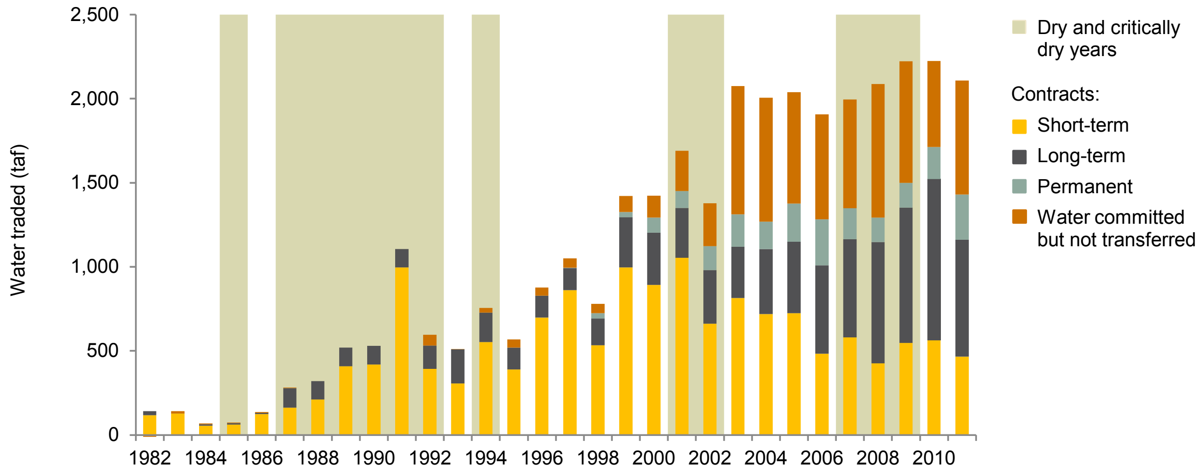

Source: Hanak and Stryjewski (2012)

between agriculture and cities remains politically sensitive (Grafton et al. 2011*b*).

4.6 Conclusions

Under standard 'release sharing' water property rights, storage and extraction decisions are made by central agencies. Water releases are allocated proportionally, subject to a small number of priority classes. User allocations can then be traded on a spot market. In the unlikely case that storage and extraction decisions are optimal and the spot market is perfect, this approach yields an efficient allocation of water.

Unsurprisingly, the economic literature on water generally advocates market trade. However, even the earliest studies are careful to acknowledge the limitations of markets, particularly externality problems. The literature puts much emphasis on spatial externalities, such as return flows. In this thesis, we focus on storage (chapter 5) and in-stream flow (chapter 7) externalities.

The above discussion — and much of this thesis — follows the conventions of the property rights school. That is, property rights are seen as the solution to externality problems: at least subject to transaction costs. Institutional Economics (IE) represents something of an alternative view. Here, emphasis is placed on the role of collectives and other institutional complexities, with or without accompanying economic theory.

This chapter provided a brief introduction to the institutional literature. Our main conclusion is that, at least with water, most of the observed complexity can be accounted for by the concept of property rights hierarchies. In keeping with the philosophy of New IE, these hierarchies can be viewed as a product of transaction costs and information problems.

This idea of property rights hierarchies may seem a touch semantic, but it illustrates an important point: that the design of water institutions is fundamentally about externalities and property rights. Whether we are interested in informal water sharing among farmers or negotiation of treaties between nations, similar economic concepts can be applied.

The water property rights systems of Australia and western US illustrate these ideas. The two systems have evolved under very different historical conditions. In Australia, the physical and institutional development of rivers has been very centralised, resulting in relatively homogeneous water rights and a flat hierarchy, that somewhat ironically, eased the transition to markets (NWC 2011*c*).

The western US evolved in a much more decentralised way, leading to a complex multi-layered hierarchy. The US approach while horrifying to some — like Alfred Deakin — is inspiring to others — like Elinor Ostrom. Although the system breeds litigation and complicates trading, it provides flexibility, allowing for adaptation of property rights to local conditions.

Surprisingly given their very different histories, both regions have converged on similar water accounting rules in regulated rivers. While we can't ignore the effects of history and rent seeking, property rights often find a way to evolve in response to fundamentals as Demsetz (1967) envisaged. The central theme of this chapter is that, while a property rights view of the world is obviously incomplete, it still has much predictive power.

Chapter 5

Water storage rights: decentralising reservoir operation

5.1 Introduction

Storage rights allow water users to hold private storage reserves in public reservoirs: partially decentralising reservoir operation. User level storage rights are now common place in the MDB. Similar rights exist in some Western US rivers. Recently, a form of storage rights even emerged on the Colorado River (Hughes 2013).

Storage rights have been examined in a number of Australian studies (Dudley and Musgrave 1988, Brennan 2008*a*, Hughes and Goesch 2009*b*). These authors demonstrate the limitations of central control of storages and argue that storage rights could improve the inter-temporal allocation of water, just as trading has improved the spatial allocation.

Given the complexity of regulated rivers, storage rights are difficult to define. 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 chapter, we compare a number of approaches to 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 capacity constraint (i.e., spills) and how they reflect evaporation losses. Such a comparison has not yet been attempted, primarily because it requires a decentralised model.

In this chapter, we present a decentralised version of our regulated river model. In the 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 transitions depend 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 region.

The goal of this chapter is to address the following policy 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, how do the answers to these questions depend on the nature of the river system?

We begin by outlining the model and defining our policy scenarios. We then summarise the literature on storage rights. The chapter then offers a brief introduction to the numerical methods employed (for more detail see chapter 8). Finally, we present the results and offer some conclusions.

5.2 The model

The basis of the model is the parametric version of the planner's problem from section 3.5.3. The decentralised version contains the same water supply and demand constraints, only here water property rights are defined, facilitating both user storage decisions and a spot market.

5.2.1 The property rights framework

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 the general form

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

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

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

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 introduced 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.

The water spot market

Users receive water allocations a_{it} adjusted for marginal delivery losses

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

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

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

Trade is subject to a transfer cost $\tau > 0$. 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. Effective user demand functions $\tilde{d}_h$ are then

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

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

The storage release rule

Actual storage releases W_t are the sum of user withdrawals plus fixed delivery losses

$$W_t = \begin{cases} \sum_{i=1}^{n} w_{it} + \delta_a/(1-\delta_b) \text{ if} & S_t > \delta_a/(1-\delta_b) \\ 0 & \text{otherwise} \end{cases}$$

When storage volumes S_t are insufficient to satisfy fixed losses no release is made. For more detail on these rules see appendix B.

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 current and expected market price for water and the physical constraints as detailed in the planner's problem (section 3.5.3).

5.3 Policy scenarios

Below we define a number of storage property rights systems. These scenarios all reflect aspects of storage right systems observed in practice. For a detailed discussion of storage rights in the MDB see appendix C.

5.3.1 Storage capacity rights (capacity sharing) — CS

Here each user is assigned a share of storage capacity and a share of inflow — for now we assume equal storage and inflow shares λ_i — such that

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

This scenario represents 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 rights of NSW and VIC on the Murray River are another example of the approach. Similar storage right arrangements exist in Northern NSW (Hughes et al. 2013) and in some US systems, such as the Texas Lower Rio-Grande.

As Dudley and Musgrave (1988) acknowledge, capacity sharing is subject to externalities, in particular ‘internal spills’: where a users account reaches its limit and excess inflow is forfeited to other users (see figure 5.1). Internal spills are zero ($x_{it+1} = 0$ for all i) if no accounts reach their limit or if all accounts reach their limit (in which case the storage physically spills). With only two users x_{it+1} is defined

$$x_{0t+1} = \begin{cases} \max\{\lambda_1.I_{t+1} - (\lambda_1 K - s_{jt} + w_{1t} + l_{1t}), 0\} & \text{if } Z_{t+1} = 0 \\ 0 & \text{otherwise} \end{cases}$$

With a large number of users calculating x_{it+1} is complicated since an initial round of internal spills may fill further accounts creating more internal spills and so on. In this case x_{it+1} can be calculated iteratively (see appendix 8, section B.3.1).

We expect capacity sharing to lead to slight under-storage (below optimal storage levels) as users will treat internal spills as pure losses: ignoring any value they have to others.

5.3.2 Spill forfeit rules (‘Spillable Water Accounts’) — SWA

Spill forfeit rules are a common alternative to 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 account volumes (see figure 5.2)

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

$$s^{'}_{it+1} = s_{it} - w_{it} - l_{it} + \lambda_i(I_{t+1} - Z_{t+1})$$

$$x_{it+1} = -Z_t\left(\frac{s^{'}_{it+1}}{\sum_{i=1}^{n} s^{'}_{it+1}}\right)$$

This type of approach is adopted in northern Victoria under the banner of ‘Spillable Water Accounts’ (SWA). Similar approaches also exist in the US, for example the San Louis Reservoir in California (see appendix C).

We expect spill forfeit rules to lead to slight over-storage, as they only partially internalise the costs of spills. That is, users will ignore the effect their storage volumes have on other users’ spill deductions.

Figure 5.1: Illustration of capacity sharing — CS

(a) Capacity sharing with four users

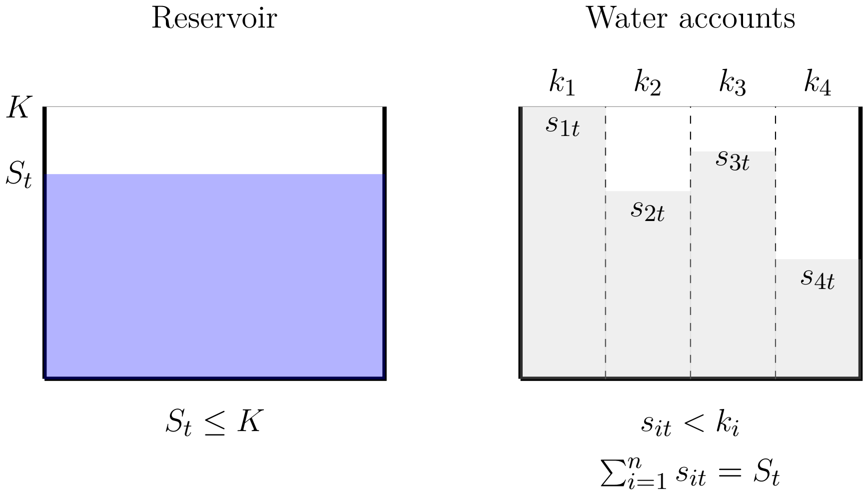

(b) Sharing of new inflow

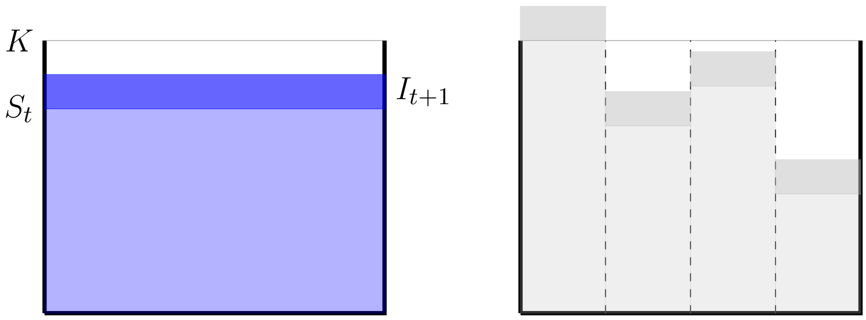

(c) Internal spills

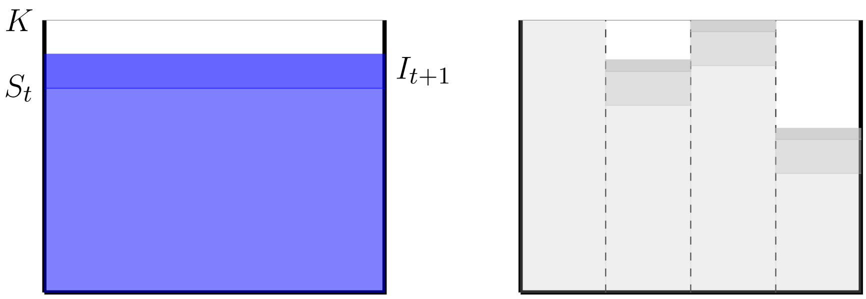

5.3.3 Capacity shares and spill forfeit rules — CS-SWA

Spill forfeit rules are often combined with storage capacity shares. Here, only water in excess of user capacity shares is subject to forfeit

$$x_{it+1} = -Z_t \left(\frac{\max\{s'_{it+1} - \lambda_i K, 0\}}{\sum_{i=1}^{n} \max\{s'_{it+1} - \lambda_i K, 0\}} \right)$$

The northern Victorian systems and the San Louis reservoir in California both involve this combination (see appendix C). Similar arrangements exist in the Macquarie river systems in NSW.

5.3.4 Open access storage (unlimited carryover) — OA

Here storage capacity is an open access resource. That is, there are no account limits and no 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 a large number of users, we expect open access to result in over-storage. While there are few examples of pure open access storage in the MDB, many systems can approach open access under certain conditions as they did in northern Victoria in 2010-11 (see appendix C).

5.3.5 No storage access ('use it or lose it') — 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 a large number of users the incentive is to 'use it or lose it': to consume or trade all water allocations — at least until 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, generally much more water was allocated than was ever used (MDBMC 1995).

Figure 5.2: Illustration of open access — OA and spill forfeit rules — SWA

(a) OA / SWA with four users

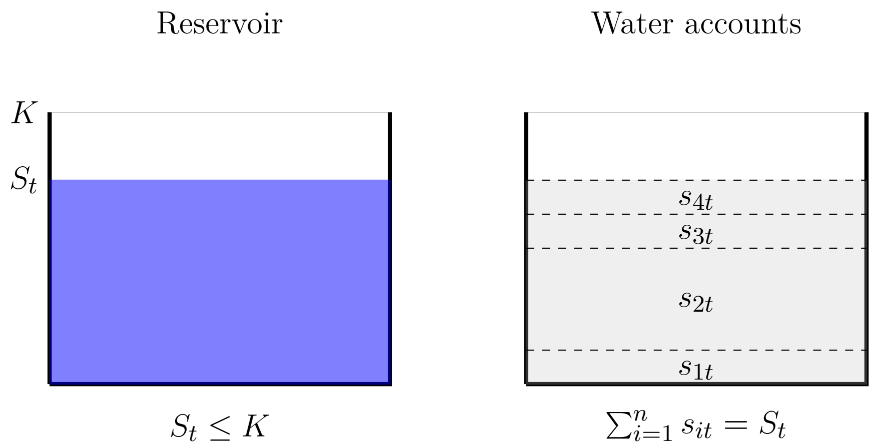

(b) OA — during a spill event

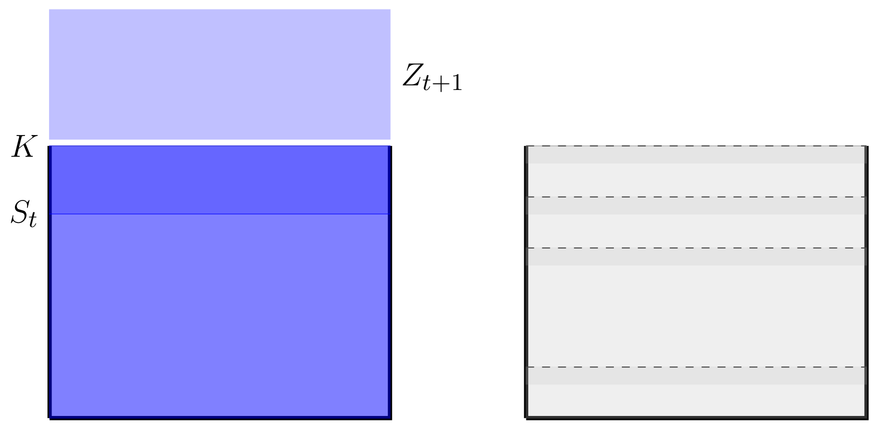

(c) SWA — effect of spill forfeit rules (relative to OA)

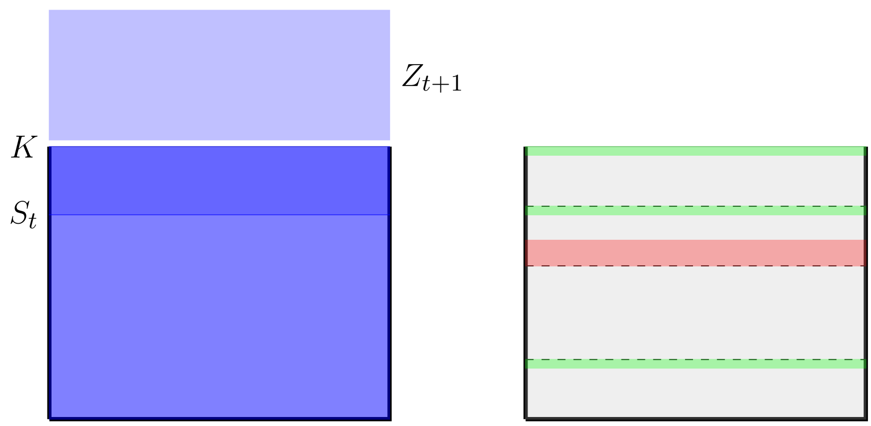

5.3.6 Storage loss deduction — CS, SWA

For the CS and SWA scenarios, we assume storage losses are allocated in proportion to user account balances

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

This is the approach adopted in southern QLD (Hughes and Goesch 2009*b*) and between NSW and VIC on the Murray. A similar approach is adopted in northern Victoria (see appendix C).

5.3.7 Socialised storage loss — CS-SL, SWA-SL

In the CS-SL and SWA-SL scenarios, we assume storage losses are allocated in proportion to inflow shares

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

This is the more common approach in the MDB (Hughes and Goesch 2009*b*).

5.4 Literature

The literature on surface water storage relies heavily on social planner models (see chapter 3)[1]. While surface water storage rights have been considered in a number of Australian studies (Dudley and Musgrave 1988, Brennan 2008*a*, Hughes and Goesch 2009*b*), they have rarely been explicitly modelled.

The literature on reservoir 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

[1] While Dudley (1988*b*) sometimes described his models as decentralised, all relied on single agent SDP methods. In some cases, Dudley (1988*b*), Dudley et al. (1998) developed models with two or three 'decision makers' (i.e., government and farmers). However, each agent was represented by an independent SDP problem such that externalities could not be taken into account.

> 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: 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)

Two sources of externality were identified: internal spills and non-linear storage losses. Dudley and Musgrave (1988) note that under two restrictive conditions: identical storage decisions and linear storage losses, the problem can be condensed to that of a representative agent. Dudley and Musgrave (1988) then present a simulation model in which users are assigned policy functions derived from the planner's solution.

Alaouze (1991) examines capacity sharing using a simplified analytical model, with no spot market, no internal spills and linear losses. Alaouze (1991) demonstrates that capacity sharing outperforms an optimal storage / proportional use allocation scenario (i.e., 'release sharing'). This result has more to do with use allocation than storage. In particular, it suggests that user storage reserves can help minimise water trade requirements (an issue we return to in chapter 6).

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*a*; 2010) evaluated central storage policy, using a model of the Goulburn region in Victoria. Brennan (2008*a*) emphasised the role of unused water allocations. Brennan (2008*a*) 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*a*; 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 'release sharing' water rights (i.e., information problems, transfer costs and policy uncertainty) and simple storage rights (i.e., externalities) relative to capacity sharing. Hughes (2010) generalises capacity sharing to more complex river systems with multiple storages, inflow sources and demand nodes. Hughes et al. (2013) detail the water storage and inflow rights systems of the MDB (see appendix C).

5.5 Solving the model

Given non-market interactions between agents, 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 practical challenges. A key contribution of this thesis, has been the development of numerical methods for efficiently solving large stochastic games. A brief outline of these techniques is provided below. For a complete treatment see chapter 8.

5.5.1 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 $\mathbf{a}_t$, 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 two policy functions $f_h(.)$ which simultaneously solve all users' problems.

Unfortunately, MPE has limited practical value in complex applied problems. Firstly, there are no general existence results and even when existence can be established, uniqueness generally can not (see chapter 8). A more practical problem with MPE is that the state space scales in the number of agents. Clearly, with large n this approach is neither feasible nor realistic.

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

In our model, 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

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

While OE is more tractable, establishing existence and uniqueness remains a problem: Weintraub et al. (2008) only establish existence conditional on MPE (see chapter 8). While some dynamic programming algorithms have been proposed around this concept, these are suited to narrow problem types where a unique equilibrium can be confirmed (see chapter 8).

5.5.2 Learning in games

Learning in games provides an alternative to the traditional focus on equilibrium. Learning models describe how players adapt in response to observed play. There is much economic literature on learning in repeated games, considering how well different learning models reflect human behaviour and if and when they 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 (Busoniu et al. 2008, Fudenberg and Levine 2007; see chapter 8).

5.5.3 A reinforcement learning approach

Reinforcement learning is a subfield of machine learning concerned with solving MDPs. Reinforcement learning algorithms optimise through simulation and so do not require an ex ante model of the 'environment' (i.e., probability transition and pay-off functions). Rather agents 'learn' optimal policies by observing the outcomes — the payoffs and state transitions — of their actions.

Our approach (detailed in chapter 8) 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 simulation is run with exploration (i.e., randomised polices) and all of the state transition, action and payoff samples are recorded. A Q or 'action-value' function is then fit to this sample. In a sense, the method translates a dynamic programming problem into a regression problem.

Under our algorithm, the Q function is optimised with respect to actions for a subsample of state points (see chapter 8). Then continuous policy f and value functions V are estimated. We call this fitted Q-V iteration. For multi-agent problems fitted Q-V iteration is combined with two repeated game type learning dynamics (i.e., smoothing rules) similar to *partial best response* and *fictitious play* (see chapter 8).

This approach provides a middle ground between the rational expectations methods (i.e., dynamic programming) of macroeconomics and the search methods (i.e., genetic algorithm) of agent based computational economics. As a simulation based learning method, it is suitable for large complex problems where the existence of an equilibrium is difficult to establish. However, computationally it remains only a small departure from rational expectation methods, ensuring our agents display near optimal (i.e., best response) behaviour.

The success of fitted Q-iteration depends crucially on function approximation, here we use a version of tile coding (Sutton and Barto 1998). For more detail on these methods see chapter 8.

5.5.4 Parameterisation

The parameterisation is identical to the planners' problem except for the transaction cost τ and the user inflow shares λ_i. For τ we assume a uniform distribution over the range \$10 to \$100 per ML.

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

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

In chapter 6 (and appendix D) we develop a model for predicting the optimal (social welfare maximising) inflow share Λ_{high} conditional on the parameter values (the number of high reliability users etc.)[2]. In this chapter, we specify a distribution for Λ_{high} relative to our estimate of the optimal share $\hat{\Lambda}_{high}$

$$\Lambda_{high} = N_0^1(\hat{\Lambda}_{high}, 0.025)$$

Inflow shares are considered in more detail in chapter 6.

5.6 Results

5.6.1 Central case

The learning algorithm

To begin, we show how the sample means of key variables evolve during the learning algorithm. Figure 5.3 shows the mean storage level $\frac{1}{T}\sum_{t=1}^{T} S_t$ at each iteration — beginning with the planner's solution as iteration 0. Each stage of the algorithm involves a simulation of around 10,000 periods. Following this, we simulate the final user policies for 500,000 periods. Our final results (section 5.6.1) are sample statistics from this last simulation.

[2]This model is based on the CS scenario. In testing, the optimal shares of the SWA, NS and OA were found to be similar.

Figure 5.3: Mean storage by iteration

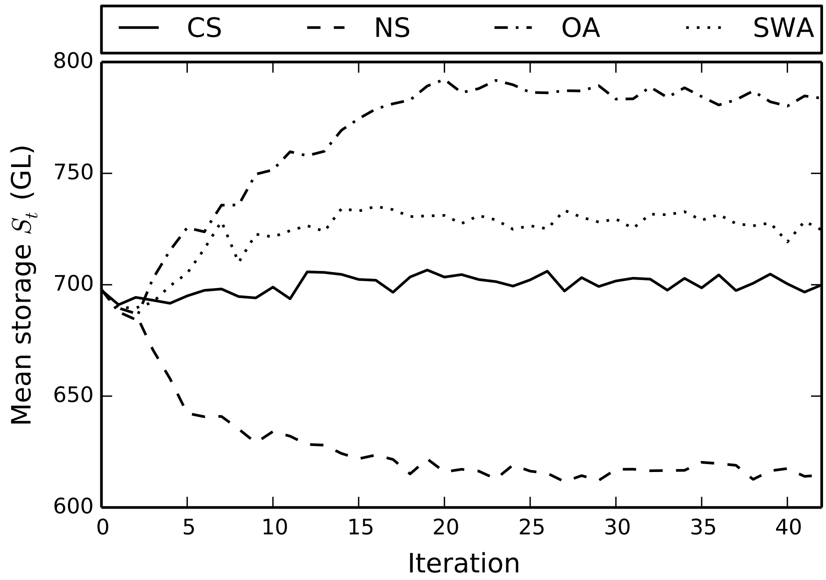

Figure 5.4: Mean social welfare by iteration

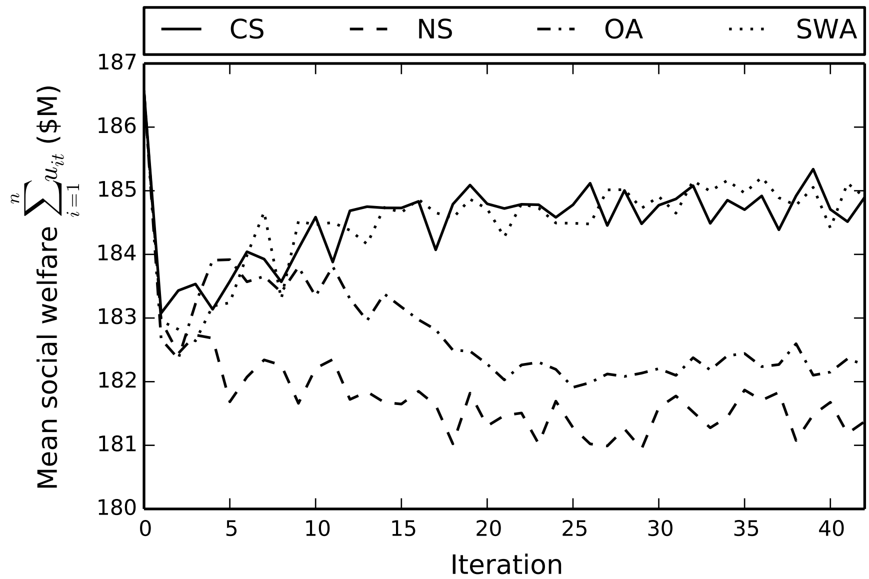

Already, some expected results are apparent. OA results in significant over-storage and NS in significant under-storage. SWA leads to slightly higher storage levels than capacity sharing. CS and SWA achieve welfare slightly below the planner's outcome, while OA and NS result in significant welfare losses.

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 5.5 and 5.6).

Figure 5.5: Value function error (Mean absolute percentage deviation)

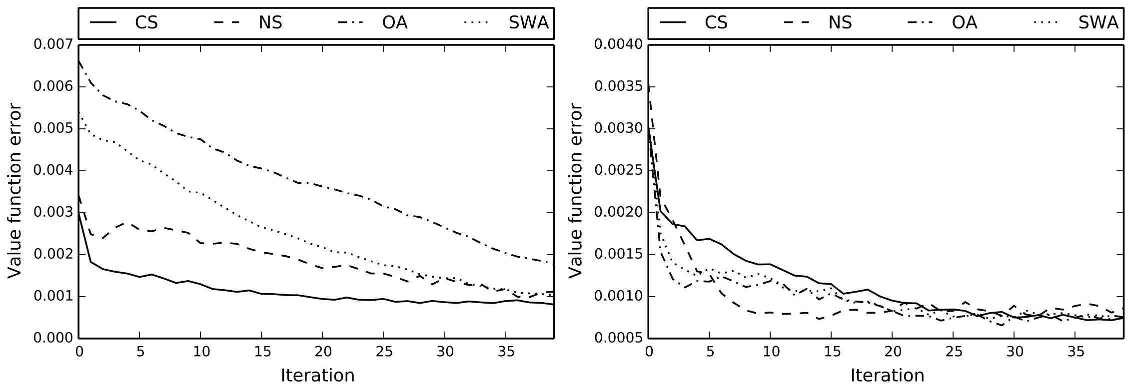

(a) Low reliability

(b) High reliability

Figure 5.6: Policy function error (Mean percentage deviation)

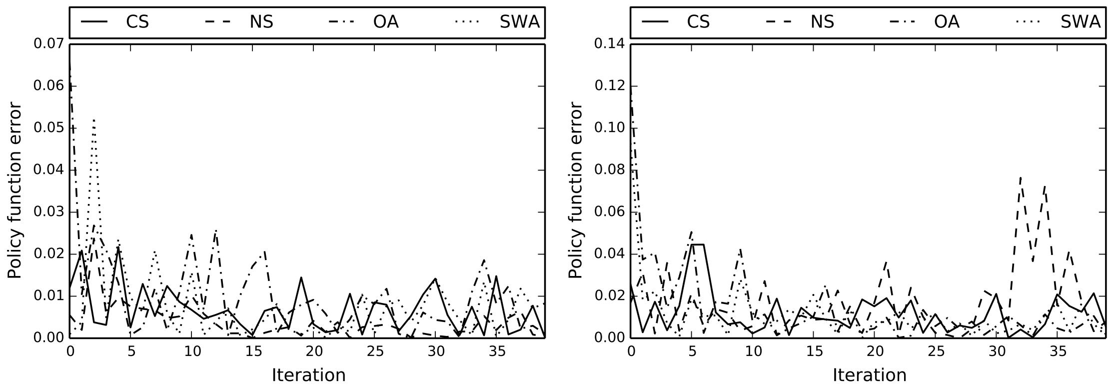

(a) Low reliability

(b) High reliability

User behaviour

Figure 5.7 shows mean aggregate withdrawals W_t conditional on the aggregate storage level. Here, we see how NS and OA induce behaviour that departs significantly from the planner's policy. Note that, mean withdrawals can be lower in spill years (when $S_t = K$) as demand for water is depressed by high inflows.

Figure 5.7: Aggregate withdrawal policy, $E[W_t|S_t]$

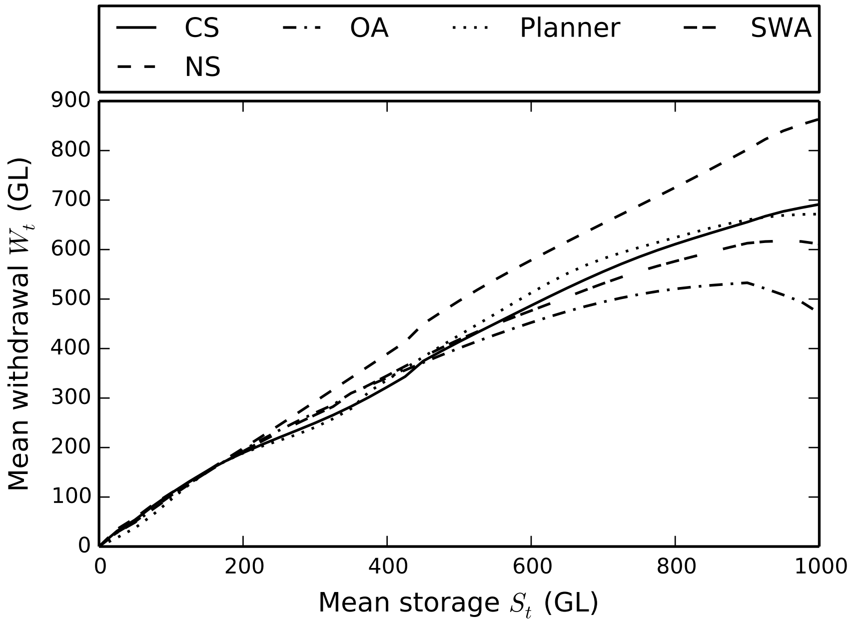

Under all scenarios, heterogeneity is observed in user storage policies. Figure 5.8 shows mean withdrawals as a proportion of account levels for high and low reliability user groups. Even with a relatively moderate transaction cost, we see specialisation: high reliability users adopt more conservative storage policy.

Figure 5.8: Mean user withdrawals w_{it} over account balances s_{it}

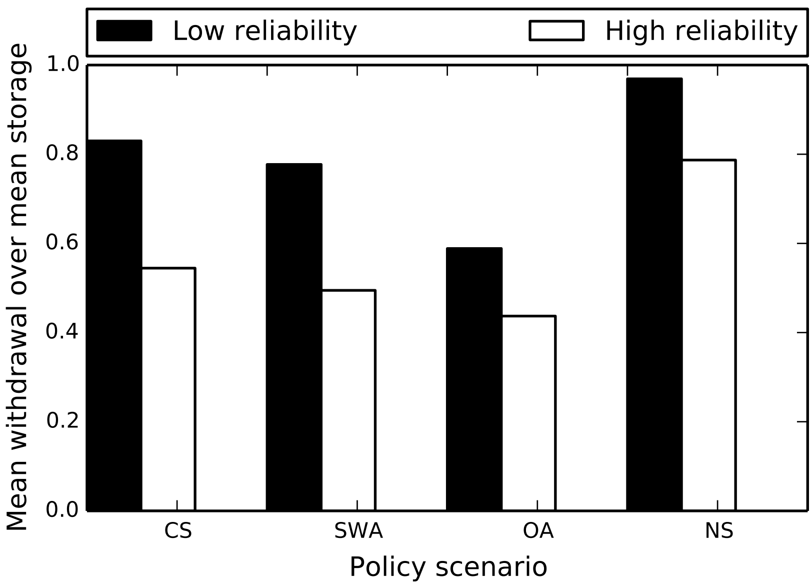

Figure 5.9 shows mean externalities x_{it} as a proportion of mean account balances

for low and high user groups. Under CS x_{it} reflects internal spills, which at least in this central case are relatively infrequent. When internal spills occur, they tend to flow from high to low reliability user accounts. Under OS, NS and SWA x_{it} reflect negative deductions for spills. These adjustments are largest under OA given the higher frequency of spills.

Figure 5.9: Mean user externalities x_{it} over account balances s_{it}

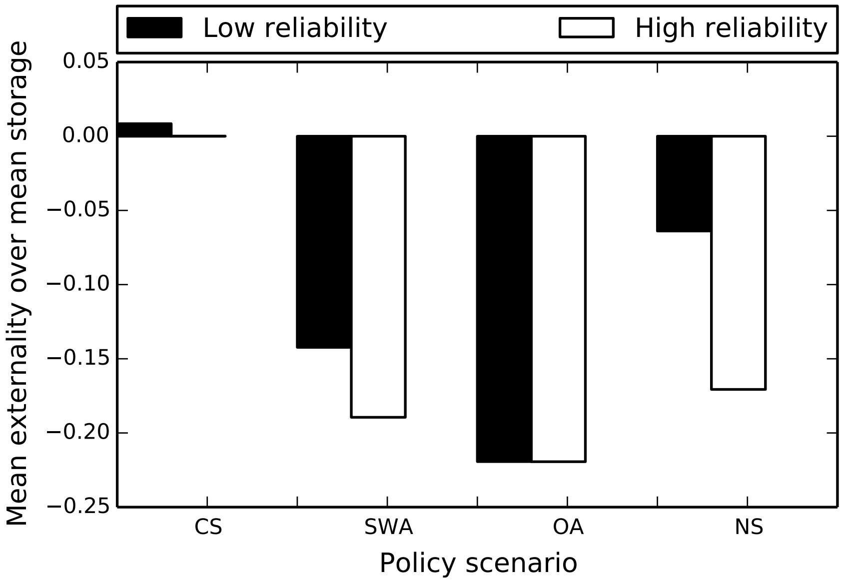

A more intuitive understanding of the model can be obtained from the simulated time series results. Figure 5.10 summarises a 50 period random sequence of results for the CS scenario. User account balances, withdrawals and net trade positions are depicted for $t = 2$ ('a wet year') and $t = 4$ ('a dry year').

These figures show high reliability users maintaining higher account balances than low users, but still tending to be net buyers during dry periods. More simulation results and further explanation is provided in appendix C.

Figure 5.10: Sample simulation results for the CS scenario

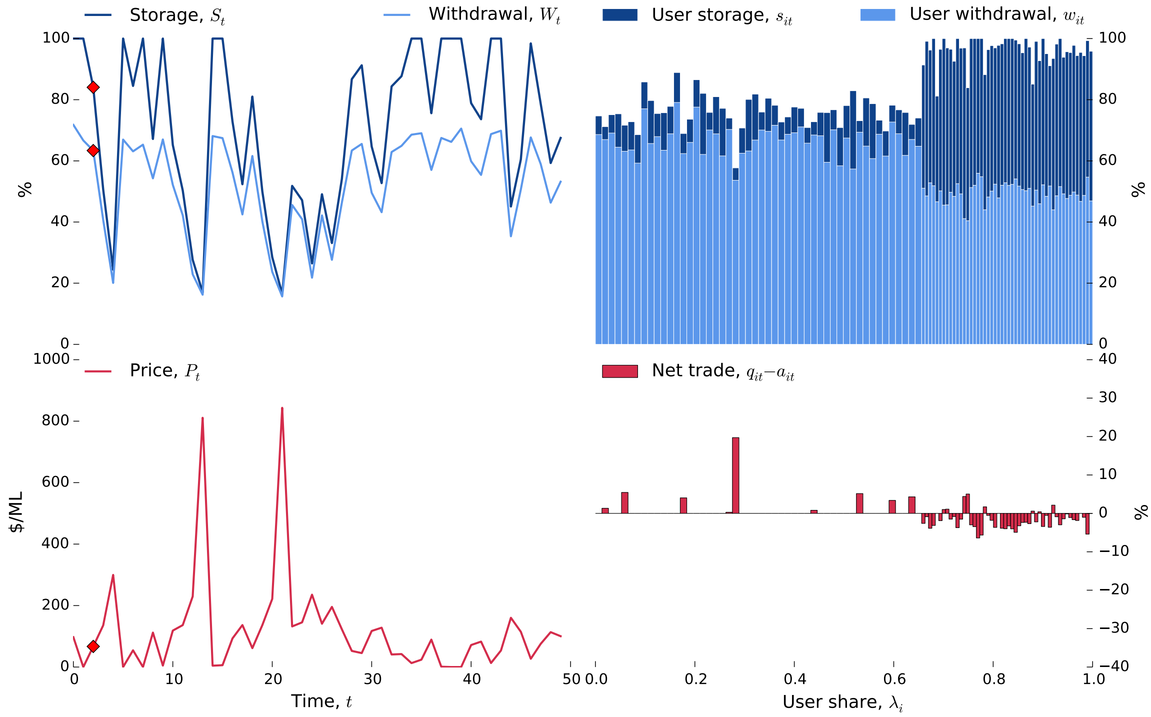

(a) $t = 2$

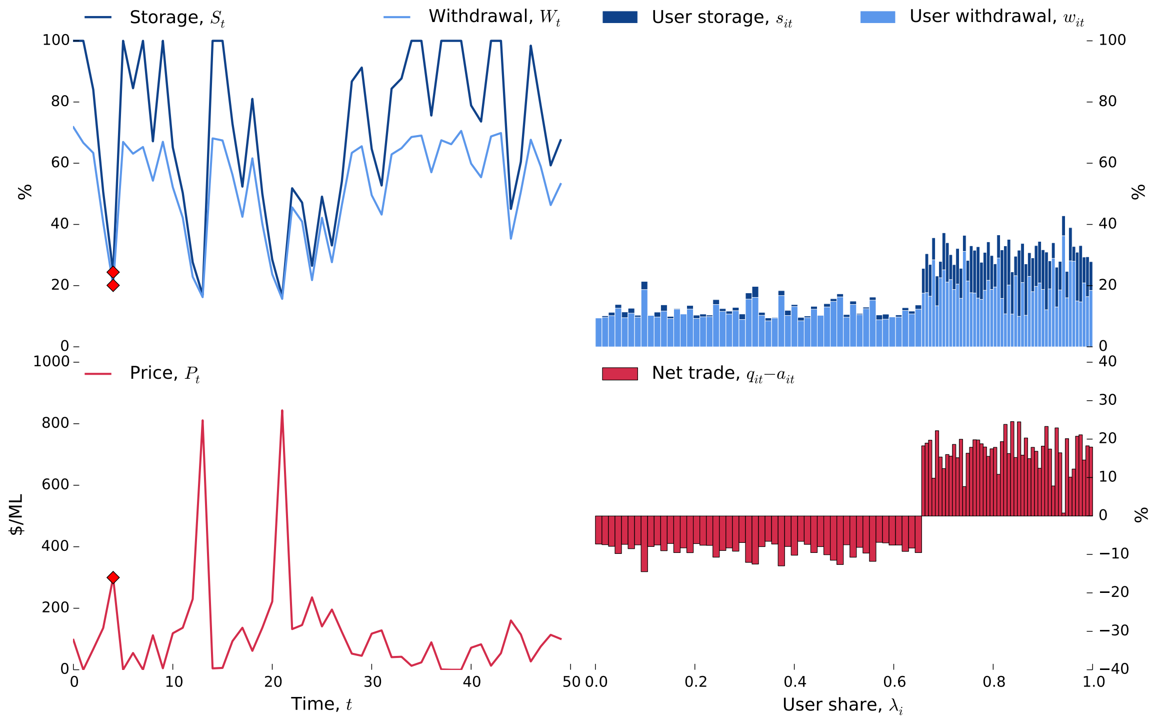

(b) $t = 4$

Final results

The results of the final simulation are summarised in tables 5.1 to 5.6.

In the central case, CS and SWA yield virtually identical mean social welfare: $184.9m slightly below the planner's outcome at $186.6m (table 5.1). OA and NS both induce welfare losses with OA ($182.3m) marginally outperforming NS ($181.4m).

Not surprisingly the mean welfare differences are relatively small (table 5.1) — at least in this central case.

As we found in chapter 3, changes in the variance of welfare are larger than changes in the mean (table 5.1). NS leads to a significant increase in the variability of withdrawals, welfare and prices (and OA to significant decreases) relative to CS and SWA.

Further, these aggregate welfare effects hide some larger distributional results. Scenarios that result in over-storage (i.e., OA) favour high reliability users at the expense of low reliability users — and vice versa (see tables 5.5 and 5.6).

While the mean welfare effects are modest, the storage differences are dramatic. Changes in storage volumes lead to large changes in the frequency and size of spill events (see table 5.4). Mean spills are 204 GL under OA, compared with 128.7 under CS and 93.1 under NS. In practice, large changes in spills will have important implications for in-stream demands like environmental flows and flood mitigation (an issue we return to in chapter 7).

Figure 5.11: Mean welfare against mean storage relative to a planner outcome

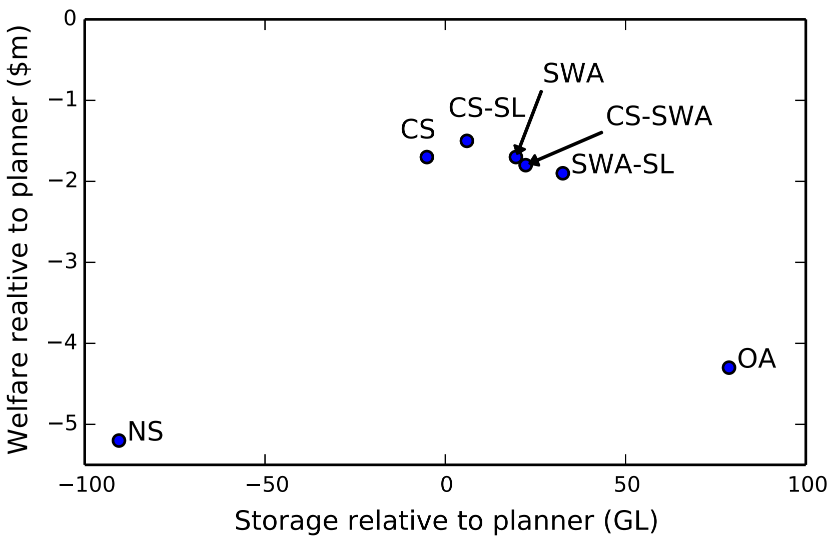

Tables 5.1 to 5.6 also contain results for scenarios CS-SL, SWA-SL and CS-SWA . As expected, socialising losses leads to increases in mean storage. At a distributional level, socialised losses tend to favour high reliability users.

The relative performance of CS, SWA, CS-SL, SWA-SL and CS-SWA can be explained by how their mean storage levels compare with optimal (figure 5.11). Note that the appropriate benchmark here is not the planner's solution, but scenario RS-O from chapter 6[3]

In this context, CS leads to below optimal storage and SWA to above optimal storage. The worst performing scenario of the group is SWA-SL as it induces significant over-storage. CS-SL is actually the preferred scenario: because the socialisation of losses offsets internal spill effects leading to near optimal storage. This is a perfect example of the second-best nature of water rights. Clearly this result will not always hold: especially for higher evaporation rates.

Finally, CS-SWA achieves storage levels slightly above (and welfare levels slightly below) SWA.

Table 5.1: Social welfare, $\sum_{i=1}^{n} u_{it}$ (\$m)

	Mean	SD	2.5th	25th	75th	97.5th
CS	184.9	25.4	123.6	176.3	201.0	208.9
CS-SL	185.1	24.6	125.6	177.7	200.1	207.9
SWA	184.9	24.5	120.6	180.1	199.1	206.7
SWA-SL	184.7	23.2	128.2	180.0	198.1	205.9
OA	182.3	20.5	134.1	177.5	193.7	201.1
NS	181.4	33.8	78.0	169.5	202.8	210.9
CS-SWA	184.8	24.6	120.3	180.1	198.9	206.6
Planner	186.6	25.8	128.2	179.0	202.2	209.4

Table 5.2: Storage, S_t (GL)

	Mean	SD	2.5th	25th	75th	97.5th
CS	699.9	278.1	163.2	471.1	1,000.0	1,000.0
CS-SL	711.0	275.4	165.6	488.2	1,000.0	1,000.0
SWA	724.6	278.0	157.0	502.6	1,000.0	1,000.0
SWA-SL	737.6	273.1	167.4	523.6	1,000.0	1,000.0
OA	783.7	263.4	180.0	600.4	1,000.0	1,000.0
NS	614.5	299.0	108.1	357.6	946.5	1,000.0
CS-SWA	727.3	279.2	156.5	503.6	1,000.0	1,000.0
Planner	697.5	282.6	159.7	462.0	1,000.0	1,000.0

[3]RS-O assumes optimal storage policy conditional on a transfer cost in the spot market and yields a mean storage level of 705 GL.

Table 5.3: Price, P_t ($ / ML)

	Mean	SD	2.5th	25th	75th	97.5th
CS	134.0	259.9	0.0	21.1	149.0	849.8
CS-SL	136.1	253.0	0.0	44.1	146.7	819.3
SWA	140.3	268.8	0.0	52.7	137.1	905.5
SWA-SL	138.0	251.6	13.9	54.2	139.2	767.4
OA	153.6	231.9	88.0	97.7	133.2	613.5
NS	182.2	390.6	0.0	0.0	167.1	1,499.8
CS-SWA	139.6	273.1	4.3	45.8	140.3	906.9
Planner	168.6	296.9	30.0	65.7	168.2	886.9

Table 5.4: Spills, Z_t (GL)

	Mean	SD	2.5th	25th	75th	97.5th
CS	128.7	277.8	0.0	0.0	82.6	1,064.5
CS-SL	134.2	283.5	0.0	0.0	100.9	1,076.2
SWA	147.3	297.4	0.0	0.0	142.1	1,110.2
SWA-SL	155.1	305.3	0.0	0.0	163.9	1,135.4
OA	204.0	347.5	0.0	0.0	288.7	1,262.1
NS	93.1	233.3	0.0	0.0	0.0	973.9
CS-SWA	151.1	301.4	0.0	0.0	152.2	1,114.3
Planner	130.4	279.8	0.0	0.0	87.9	1,054.7

Table 5.5: Total low reliability payoff, $\sum_{i \in \mathcal{U}^{low}} u_{it}$ ($m)

	Mean	SD	2.5th	25th	75th	97.5th
CS	80.8	14.4	51.1	69.4	92.6	97.0
CS-SL	80.6	14.0	50.9	70.1	91.9	96.3
SWA	80.2	13.0	50.6	71.6	90.1	94.4
SWA-SL	79.5	13.1	48.6	71.6	89.2	93.6
OA	77.2	9.6	51.8	74.8	83.4	89.2
NS	82.3	14.1	55.3	71.3	94.7	99.7
CS-SWA	80.3	12.6	51.5	72.5	89.5	93.9
Planner	79.2	18.6	40.5	68.2	92.8	96.9

Table 5.6: Total high reliability payoff, $\sum_{i \in \mathcal{U}^{high}} u_{it}$ (\$m)

	Mean	SD	2.5th	25th	75th	97.5th
CS	104.1	15.6	56.6	102.1	111.2	115.5
CS-SL	104.5	15.0	60.2	102.7	111.2	115.4
SWA	104.7	16.4	51.2	103.6	111.5	115.6
SWA-SL	105.2	14.9	65.1	103.9	111.3	115.5
OA	105.1	14.3	71.9	103.8	110.9	114.9
NS	99.1	24.6	8.8	99.5	111.0	115.7
CS-SWA	104.5	16.7	50.1	103.3	111.4	115.5
Planner	107.3	13.7	87.3	106.7	112.2	115.9

5.6.2 General case

Here we present the results of 1,500 model runs. For each run we draw a parameter set randomly from our defined distributions, then solve the CS, SWA, OA and NS scenarios. For each parameter set and each scenario, we calculate the following sample means:

- Social welfare: $\frac{1}{T}\sum_{t=1}^{T}\sum_{i=1}^{n} u_{it}$
- Low and high reliability welfare: $\frac{1}{T}\sum_{t=1}^{T}\sum_{i \in \mathcal{U}^{low}} u_{it}$, $\frac{1}{T}\sum_{t=1}^{T}\sum_{i \in \mathcal{U}^{high}} u_{it}$
- Storage: $\frac{1}{T}\sum_{t=1}^{T} S_{it}$
- Spills: $\frac{1}{T}\sum_{t=1}^{T} Z_{it}$

For each statistic and each model run we also define an index relative to the CS scenario. Summary statistics are presented at the end of this section in tables 5.7 to 5.18.

Social welfare

Social welfare results are summarised in tables 5.7 and 5.8 and figure 5.12.

Figure 5.12: Social welfare index

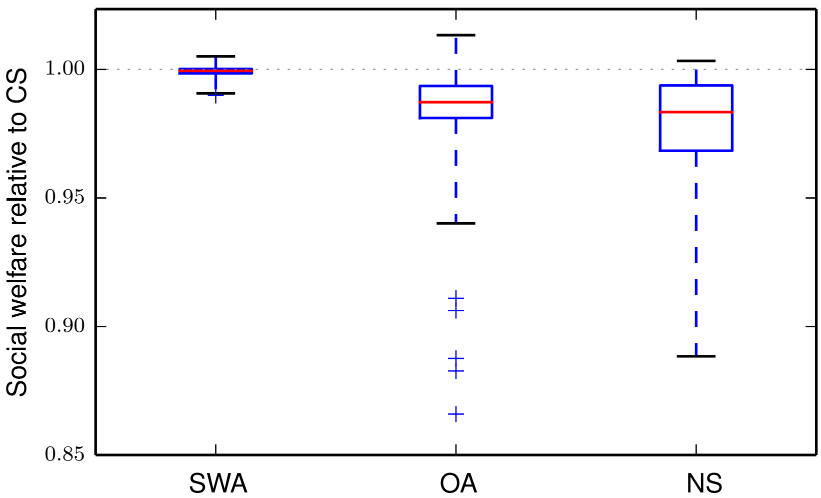

On average, CS achieves the highest mean welfare. On social welfare grounds, CS is the preferred scenario in 841 cases, SWA in 409, NS in 115 and OA in 49[4]. In those cases where CS is not preferred, the welfare differences are small. The welfare differences between SWA and CS are almost always trivial, but large welfare losses are sometimes observed with NS and OA.

Next we look at the correlation between the welfare results and the model parameters. First, the mean welfare indexes are regressed against the parameters (using the method of random forests, see appendix B).

The two most important parameters are: the ratio of mean inflow to storage capacity and inflow variation (table 5.10). In low inflow (and high variance) cases NS performs poorly (figure 5.13). An inverse result is observed for the OA scenario: it performs poorly in rivers with high inflow relative to capacity (i.e., more frequent spills).

[4] 51 runs were excluded due to numerical errors and in 35 cases the results were too close to call.

Figure 5.13: Welfare index regression results

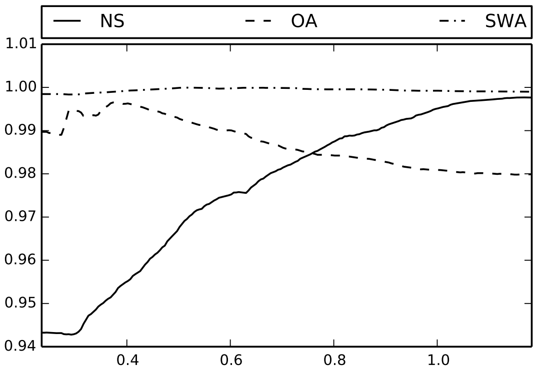

(a) Mean inflow to capacity, $E[I_t]/K$

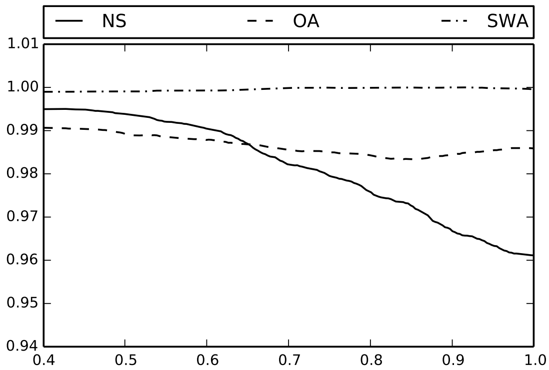

(b) Inflow variability, cv_I

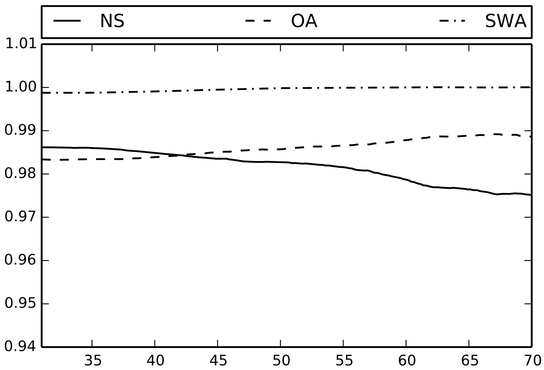

(c) Number of high reliability users, n_{high}

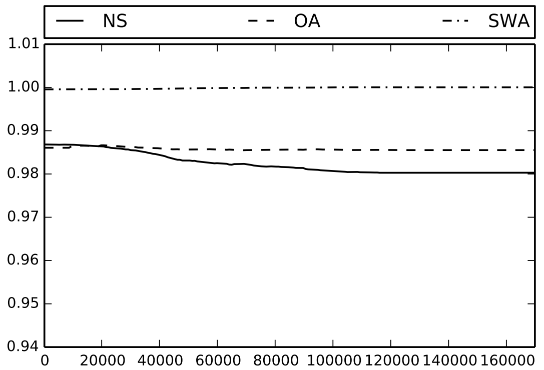

(d) Fixed delivery loss, δ_{1a}

Figure 5.14 shows the preferred scenario (CS red dots, SWA yellow dots, OA blue dots and NS green dots) against the two inflow parameters for the 1,500 model runs. Here we see that NS is preferred only in low mean inflow and high inflow variance cases (green dots, lower right), while OA is preferred only in low inflow cases (blue dots, left).

Figure 5.14: Preferred scenario by inflow parameters

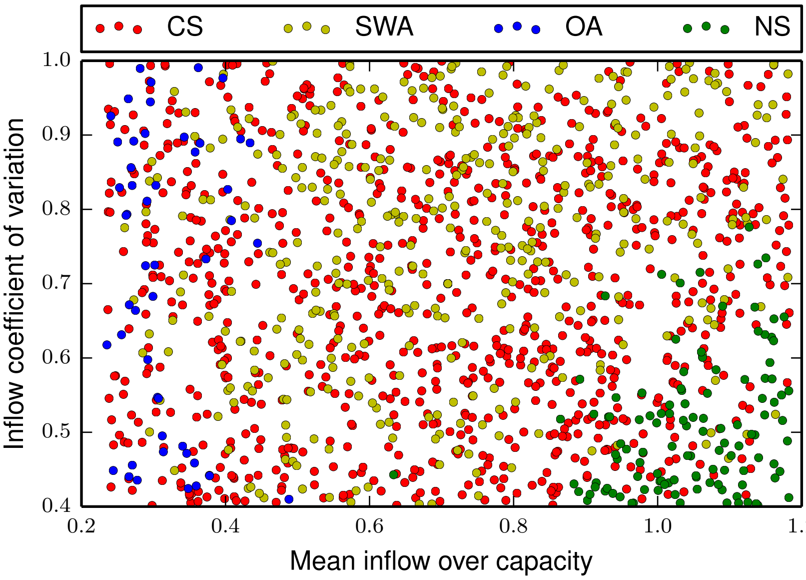

Figure 5.15: Preferred scenario by inflow parameters, with classifier predictions

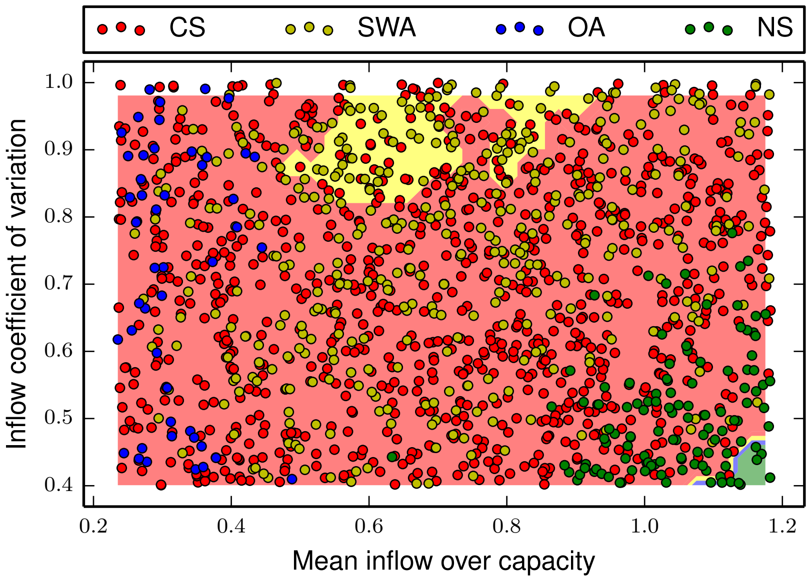

Next the preferred scenario was regressed — as a qualitative dependent variable — against the parameters (using a random forest classifier). The shaded areas in figure 5.15 represent the regression model's prediction of the preferred scenario (given other parameters fixed at sample means). The regression model predicts CS as the best scenario for the vast majority of the parameter space. SWA is preferred in some high variance cases (the shaded yellow area).

Figure 5.16 shows the regression model predictions for other slices of the parameter space. OA is more likely to be preferred in rivers with more high reliability users (figure 5.16a) and NS in cases with fewer high reliability users (figure 5.16b). SWA is more often preferred in markets with low transaction costs (figure 5.16d).

Figure 5.16: Predicted preferred scenario by inflow parameters

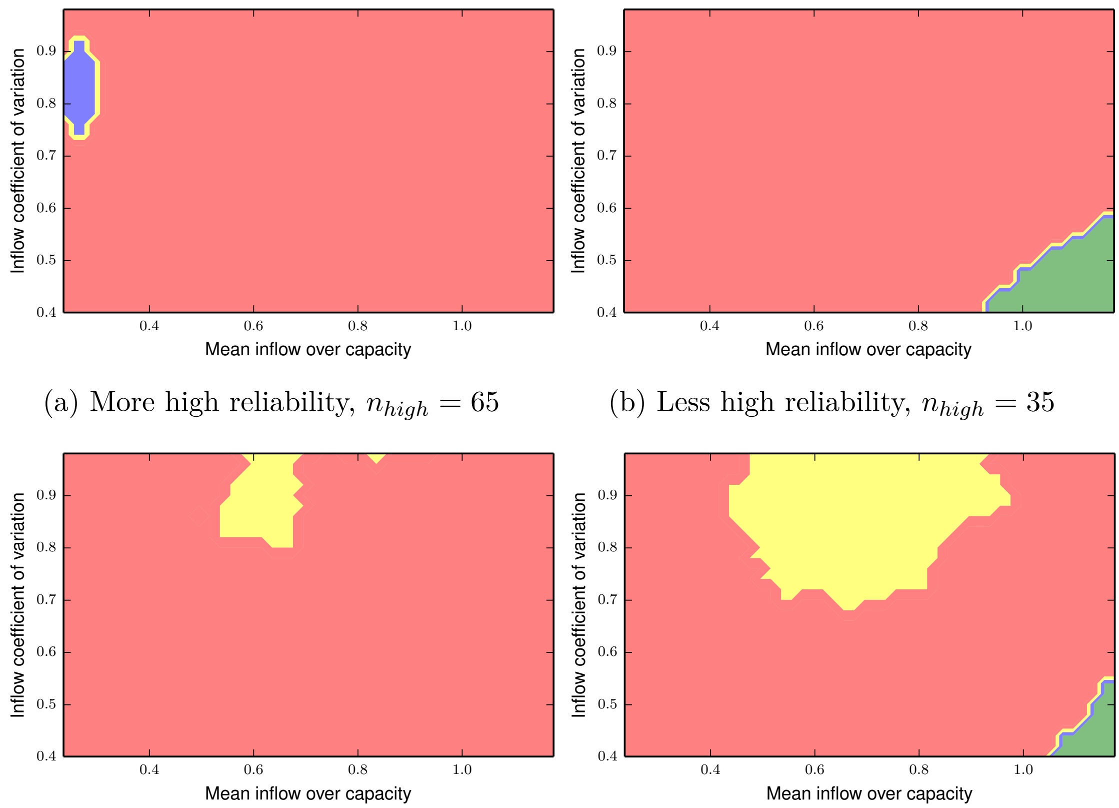

(a) More high reliability, $n_{high} = 65$

(b) Less high reliability, $n_{high} = 35$

(c) High transaction cost, $\tau = 100$

(d) Low transaction cost, $\tau = 20$

Welfare distribution

The distributional effects are generally larger than the mean welfare effects, see tables 5.11 to 5.14 and figures 5.17 and 5.18 below. In general, OA favours high reliability users and NS low. Although SWA and CS are barely separable on mean welfare grounds, there are some noticeable distributional effects, with SWA favouring high reliability users.

Figure 5.17: Low reliability payoff index

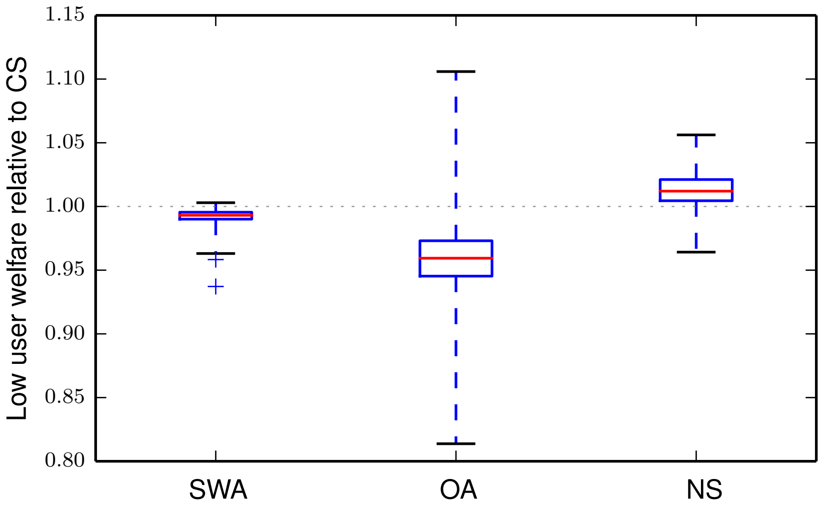

Figure 5.18: High reliability payoff index

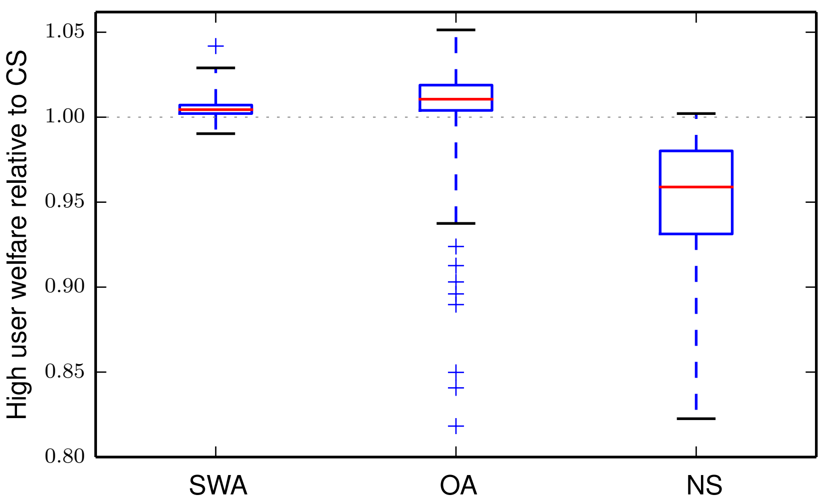

Storage

The scenarios all induce significant changes in mean storage levels (tables 5.15 and 5.16 and figure 5.19). In almost all cases, OA induces significantly higher storage reserves than CS, NS significantly lower and SWA slightly higher.

Figure 5.19: Storage index

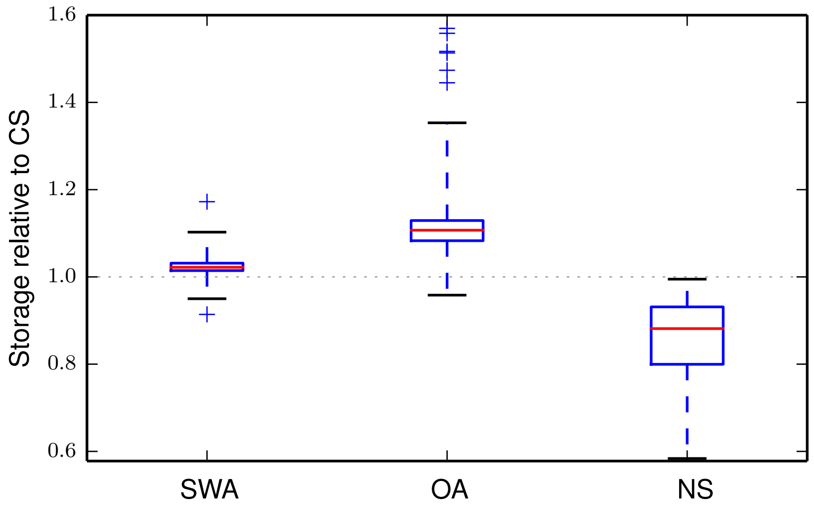

Spills

Changes in mean storage levels lead to amplified changes in spills (tables 5.17 and 5.18 and figure 5.20). Higher mean spills reflect both an increase in the frequency and magnitude of spill events.

Figure 5.20: Spills index

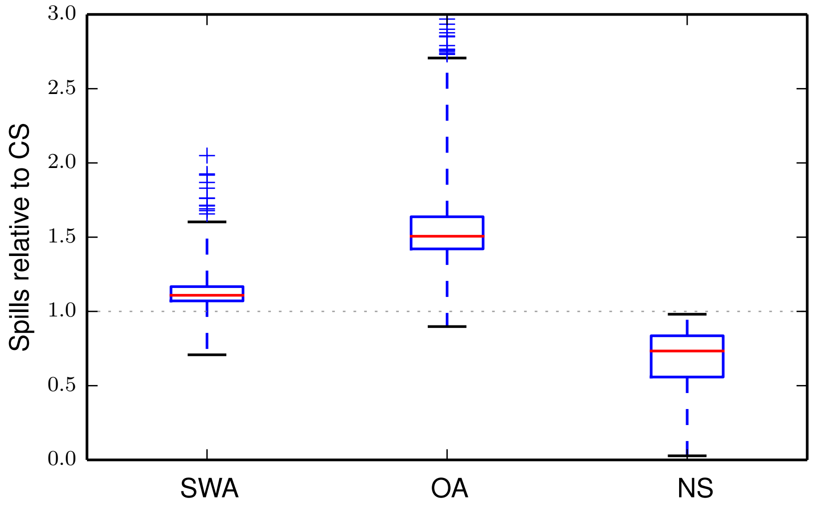

Table 5.7: Mean social welfare, $\frac{1}{T}\sum_{t=1}^{T}\sum_{i=1}^{n} u_{it}$ ($m)

	Mean	Min	Q1	Q3	Max
CS	186.68	44.60	132.48	233.51	408.34
SWA	186.54	44.49	132.28	233.49	408.27
OA	183.83	44.50	131.66	229.51	402.84
NS	183.55	41.69	129.20	231.84	408.46
Planner	188.84	45.03	134.43	236.08	418.30

Table 5.8: Social welfare index

	Mean	Median	Min	Q1	Q3	Max
CS	1.0000	1.0000	1.0000	1.0000	1.0000	1.0000
SWA	0.9992	0.9994	0.9899	0.9985	1.0002	1.0051
OA	0.9865	0.9873	0.8057	0.9811	0.9936	1.0133
NS	0.9784	0.9834	0.8884	0.9684	0.9938	1.0033

Table 5.9: Preferred scenario classifier: parameter importance and sample means

	Importance	CS	SWA	OA	NS
$E[I]/K$	11.65	0.71	0.72	0.32	1.05
c_v	10.43	0.69	0.75	0.71	0.51
n_{high}	9.39	48.98	53.46	56.90	43.30
$\Lambda_{high} - \hat{\Lambda}_{high}$	8.87	0.00	-0.00	0.00	-0.00
δ_a	7.35	52,022.27	60,278.51	34,108.12	62,847.92
τ	7.06	55.74	51.96	53.64	50.51
$\frac{\mathscr{A}_{low}}{E[I]/K}$	6.70	6,807.20	6,912.56	7,199.59	7,091.53
α	6.69	8.89	9.58	9.58	9.04
δ_0	6.67	0.62	0.63	0.64	0.61
σ_η	6.37	0.15	0.15	0.15	0.15
ρ_e	6.35	0.40	0.39	0.40	0.41
δ_b	6.31	0.23	0.23	0.22	0.23
ρ_I	6.17	0.25	0.25	0.25	0.25

Table 5.10: Welfare index regression: parameter importance

	Importance
$E[I]/K$	39.31
c_v	26.87
δ_{1a}	7.86
n_{high}	7.42
$\frac{\mathscr{A}_{low}}{E[I]/K}$	3.21
α	2.61
$\delta 0$	2.14
τ	1.99
ρ_I	1.84
δ_{1b}	1.83
ρ_e	1.68
$\Lambda_{high} - \hat{\Lambda}_{high}$	1.65
σ_η	1.58

Table 5.11: 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	83.07	21.59	55.44	105.53	182.71
SWA	82.40	21.46	54.90	104.65	181.67
OA	79.41	20.88	53.29	100.80	174.61
NS	84.20	21.55	55.76	107.31	183.93

Table 5.12: Low reliability payoff index

	Mean	Min	Q1	Q3	Max
CS	1.000	1.000	1.000	1.000	1.000
SWA	0.992	0.937	0.990	0.995	1.003
OA	0.959	0.747	0.945	0.973	1.199
NS	1.012	0.964	1.004	1.021	1.056

Table 5.13: 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	103.62	15.76	70.60	130.73	294.22
SWA	104.13	15.75	70.96	131.28	297.28
OA	104.42	16.15	71.31	131.54	296.48
NS	99.34	13.53	65.38	126.88	288.04

Table 5.14: High reliability payoff index

	Mean	Min	Q1	Q3	Max
CS	1.000	1.000	1.000	1.000	1.000
SWA	1.005	0.990	1.002	1.007	1.042
OA	1.010	0.768	1.004	1.019	1.051
NS	0.952	0.823	0.931	0.980	1.002

Table 5.15: Mean storage, $\frac{1}{T}\sum_{t=1}^{T} S_t$ (GL)

	Mean	Min	Q1	Q3	Max
CS	684.16	302.69	625.56	758.55	908.17
SWA	699.71	291.70	644.70	774.49	925.24
OA	757.35	329.24	706.39	835.53	999.85
NS	596.03	229.72	502.18	705.46	892.67
Planner	684.37	326.86	633.25	753.90	892.06

Table 5.16: Storage index

	Mean	Min	Q1	Q3	Max
CS	1.000	1.000	1.000	1.000	1.000
SWA	1.023	0.914	1.014	1.031	1.172
OA	1.110	0.958	1.083	1.129	1.852
NS	0.858	0.584	0.800	0.931	0.995

Table 5.17: Mean spills, $\frac{1}{T}\sum_{t=1}^{T} Z_t$ (GL)

	Mean	Min	Q1	Q3	Max
CS	134.43	0.14	59.45	200.26	349.92
SWA	147.26	0.10	70.70	217.72	358.37
OA	198.94	0.36	98.19	290.64	517.09
NS	105.79	0.01	33.76	166.40	300.39
Planner	133.78	0.31	60.82	197.49	353.44

Table 5.18: Spills index

	Mean	Min	Q1	Q3	Max
CS	1.000	1.000	1.000	1.000	1.000
SWA	1.139	0.708	1.071	1.167	2.049
OA	1.621	0.898	1.421	1.637	15.417
NS	0.670	0.028	0.558	0.836	0.981

5.7 Conclusions

5.7.1 Research questions

We're now in a position to address our initial questions:

Which system of storage rights maximises social welfare?

The simple answer is capacity sharing (storage capacity rights). 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 mean 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 higher mean storage levels than 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*a*, 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.

5.7.2 Policy implications

In many cases the welfare differences between scenarios are trivial and will be outweighed by transition costs. In others, a transition from no storage access to some form of storage right may offer a significant gain. However, this gain can go unrealised if storage rights lead to 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 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 on the river system. Given our results, it is understandable that capacity sharing has been implemented in the northern MDB (where inflows are high relative to storage) and spill forfeit rules / open access in the south (where spills are less frequent). While approaches may be adapted to local conditions the preferred approach may change in response to new developments — particularly climate change. For example, expectations of lower and more variable inflows due to climate change help explain the recent shift from NS to storage rights in the MDB (and for that matter the Colorado).

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

Chapter 6

Water flow rights: proportional versus priority

6.1 Introduction

Priority rules are a feature of almost all water property right systems. Priority rules define distinct classes of rights with varying degrees of 'reliability'. Here reliability refers broadly to the variance of users' pre-trade water allocations[1].

As Freebairn and Quiggin (2006) note, priority rules can be justified by spot market frictions. Providing priority to users with more inelastic demands (i.e., high reliability users) can reduce trade requirements and exposure to transfer costs in comparison with proportional rights.

Priority rights have also been justified on risk management grounds (Lefebvre et al. 2012, Beare 2010). Water users more averse or exposed to water supply risks, may value priority rights for their ability to minimise variation in total income.

Proportional versus priority flow rights has been a long standing debate in water economics (see Burness and Quirk 1979). In this chapter, we reconsider this issue in the context of storage. A key question, is whether priority rules are still necessary if we have well defined storage rights.

While storage rights systems (such as capacity sharing) often involve proportional inflow shares, users can determine their allocations through storage decisions. In chapter 5 our high reliability users maintained greater reserves in order to minimise their reliance on trade — even in scenarios with relatively low transfer costs.

[1] In this context, reliability is often defined more narrowly as the proportion of years users receive their target allocation, see appendix D.

In this chapter, we put less emphasis on inter-temporal efficiency and more on use allocation. We consider a set of release sharing scenarios where a planner makes storage decisions and users receive either proportional or priority shares of releases. We compare these against a number of capacity sharing based scenarios. We do this using essentially the same model and parameter assumptions adopted in chapter 5.

We begin by defining our policy scenarios. We then review the economic literature on priority water rights. Next we introduce some minor changes to the model setup and solution methods. Finally, we summarise the results and provide conclusions.

6.2 Policy scenarios

We consider two types of policy scenario: capacity sharing and 'release sharing'. With release sharing, storage decisions are made by a central planner and users receive shares in releases (as with our standard water rights from chapter 4). For now we adopt the optimistic assumption that our planner has full information and is a social welfare maximiser. Under capacity sharing users receive shares in inflow and make their own storage decisions (as they did in chapter 5).

In both cases, storage decisions will be near optimal. However, the scenarios differ in how the release / inflow shares are specified: being either proportional or priority based.

Since we are focused on use allocation here, the user shares λ_i (i.e., the endowments) become important. We discuss the setting of these shares at the end of the section.

6.2.1 Release sharing — proportional — RS

Here a planner determines W_t in order to maximise social welfare. Users receive proportional shares in allocations A_t

$$A_t = \max\{W_t(1-\delta_b) - \delta_a, 0\}$$

$$a_{it} = \lambda_i A_t$$

$$\sum_{i=1}^{n} \lambda_i = 1$$

6.2.2 Release sharing — two priority classes — RS-HL

Once again releases are determined centrally, however here we have two priority classes, with high reliability users having priority over low

$$a_{it} = \begin{cases} \min\{\lambda_i A_t, \Lambda_{high}.\bar{A}\} & \text{if } i \in \mathcal{U}_{high} \\ \lambda_i(\max\{A_t - \Lambda_{high}\bar{A}, 0\}) & \text{if } i \in \mathcal{U}_{low} \end{cases}$$

where $\Lambda_{high} = \sum_{i \in \mathcal{U}_{high}} \lambda_i$ and $\bar{A}$ is the maximum possible allocation

$$\bar{A} = (K(1 - \delta_b) - \delta_a)$$

6.2.3 Capacity sharing — CS

Here we have the CS scenario from chapter 5

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

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

6.2.4 Capacity sharing — unbundled — CS-U

Here we allow for separate (unbundled) inflow rights λ_i^I and capacity rights λ_i^K such that

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

$$\sum_{i=1}^{n} \lambda_i^I = 1$$
$$\sum_{i=1}^{n} \lambda_i^K = 1$$

While not common in practice, unbundling of storage and inflow capacity was considered a possibility by Dudley and Musgrave (1988) (although their modeling assumes equal inflow and capacity shares).

6.2.5 Capacity sharing — priority — CS-HL

Here we include priority inflow rights under a capacity sharing system

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

$$a_{it} = \begin{cases} \frac{\lambda_i}{\Lambda_{high}} \min\{I_t, \Lambda_{high}\bar{A}\} & \text{if } i \in \mathcal{U}_{high} \\ \frac{\lambda_i}{(1-\Lambda_{high})} \max\{I_t - \Lambda_{high}\bar{A}, 0\} & \text{if } i \in \mathcal{U}_{low} \end{cases}$$

Such a combination of priority and storage capacity rights has been proposed previously by Truong et al. (2010).

6.2.6 Inflow / release shares

Since the users within each class are ex ante identical, we assume they hold identical shares

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

This generalises to scenario CS-U where we have Λ^I_{high} and Λ^K_{high}.

For each scenario we consider two cases: shares in proportion to user target demands ($\Lambda_{high} = \bar{Q}_{high}/\bar{Q}$) and optimal (social welfare maximising) shares. We denote optimal share scenarios with the suffix -O.

Optimal shares can be interpreted as the product of a water share (i.e., permanent entitlement) market or an informed planner. Note that shares are always time invariant: shares are set (the share market clears) at time $t = 0$. That is, relative transfer costs force users to take a long run position in the water share market and rely on the spot market for all short-run (annual) adjustment.

We also assume users hold only one class of right. That is, high reliability users hold only high reliability rights: there are no right 'portfolios'[2]. Relaxing this constraint remains a potential extension of this work. Note that the reliability levels of high and low rights are endogenously determined through the planner's storage policy and the choice of Λ_{high} (see appendix D).

We discuss the computation of optimal Λ_{high} in section 6.5.1. In appendix 6.4 we consider some of the problems faced in setting Λ_{high} in real world systems.

[2]Note that proportional rights are equivalent to a priority rights scenario where each user has a pro rata share of each right. In that sense we consider two potential water right portfolios: high users have all the priority rights and each user has a pro rata share.

Table 6.1: Policy scenario summary

Scenario	Storage	Capacity rights	Flow sharing	Share endowment
RS	planner	None	proportional	arbitrary
RS-O	planner	None	proportional	optimal
RS-HL	planner	None	priority	arbitrary
RS-HL-O	planner	None	priority	optimal
CS	decentralised	bundled	proportional	arbitrary
CS-O	decentralised	bundled	proportional	optimal
CS-HL	decentralised	bundled	priority	arbitrary
CS-HL-O	decentralised	bundled	priority	optimal
CS-U	decentralised	unbundled	proportional	optimal

6.3 Literature

Among the earliest theoretical work on the subject is that of Burness and Quirk (1979), who consider individual priority ordering (section 3.4.3) under prior appropriation water rights. Burness and Quirk (1979) use a simple theoretical model: an unregulated river with stochastic flow, n homogeneous risk neutral water users and no spot market. Their key result is that individual priority ordering is inefficient relative to equal (proportional) sharing — senior rights holders receive too much water and junior too little.

Burness and Quirk (1979) go on to show that with a perfect water share market, priority ordering can replicate equal sharing. However, this involves an unrealistic equilibrium where each user holds a portfolio of n entitlements, containing a $1/n$ share of each priority class.

In a working paper, Burness and Quirk (1977) consider the case of regulated rivers: assuming an optimising storage manager and again no spot market. They find similar results to the unregulated case — priority ordering is inferior to equal sharing unless there is a perfect share market.

As Burness and Quirk (1977) note, the real strength of prior appropriation was in providing investment confidence during a development period. These days, strict date ordering is rarely used to ration water in regulated rivers (see chapter 4). So for both theoretical and practical reasons, we do not consider individual priority any further.

Freebairn and Quiggin (2006) consider simple priority rules (as in our RS-HL scenario), using a theoretical model, with two water availability states, two types of water user and no storage. Freebairn and Quiggin (2006) show that proportional and priority rights are equivalent under perfect spot markets, but that priority rights are preferred under positive transfer costs. Adamson et al. (2006) provide a numerical demonstration for the MDB.

Freebairn and Quiggin (2006) go on to consider how the mix of priority rights (i.e., our Λ_{high} parameter) might be set and adjusted over time. Here, they propose a type of water share conversion market (we return to this issue in section 6.4).

Lefebvre et al. (2012) compare priority and proportional water sharing within laboratory experiments. Their experimental design includes two types of water user (low and high reliability) and allows for both trade in shares and water allocations. Lefebvre et al. (2012) consider two motivations for priority rights, transfer costs and risk aversion.

The experimental results of Lefebvre et al. (2012) provide support for both the transfer cost and risk aversion arguments. Priority rights result in an increase in average profits (at least when transfer costs in the spot market are lower than in the share market) and decreases in income volatility (at least for the priority users).

The risk preference motivation for priority rights is considered further by Beare (2010). Beare (2010) emphasises the role of water rights as risky assets, with returns that are inversely correlated with farming returns. Beare (2010) estimates a risk premium for priority water rights in the MDB using the Capital Asset Pricing Model (CAPM).

Truong et al. (2010) consider both storage rights and priority rules. Truong et al. (2010) suggest extending prioritisation to a capacity sharing system (i.e., priority inflow rights). Truong et al. (2010) note, that while priority rights may limit spot market transactions, they can complicate share market transactions.

Hughes and Goesch (2009*b*) argue that storage rights — even with proportional inflow shares — can help minimise spot market transactions. Hughes and Goesch (2009*b*) also argue that priority rights can complicate share market transactions by creating heterogeneity and more generally contributing to policy uncertainty. We provide further remarks on these issues below.

6.4 Priority rights in practice

In this section we address some aspects of priority rights not represented by our model.

The first is policy uncertainty. As discussed back in section 3.4.4, priority rights tend to concentrate policy uncertainty over water allocation on to the 'junior' priority classes: so that small changes in policy can have large effects on certain users.

The Westlands irrigation district in California provides a good example of the problem. Figure 6.1 compares water availability (storage volumes) against allocations. Westland's now receives much lower allocations for equivalent levels of supply. As a 'junior' right holder Westlands has been exposed to policy changes — in this case new environmental requirements.

Figure 6.1: Storage levels versus Westlands allocations

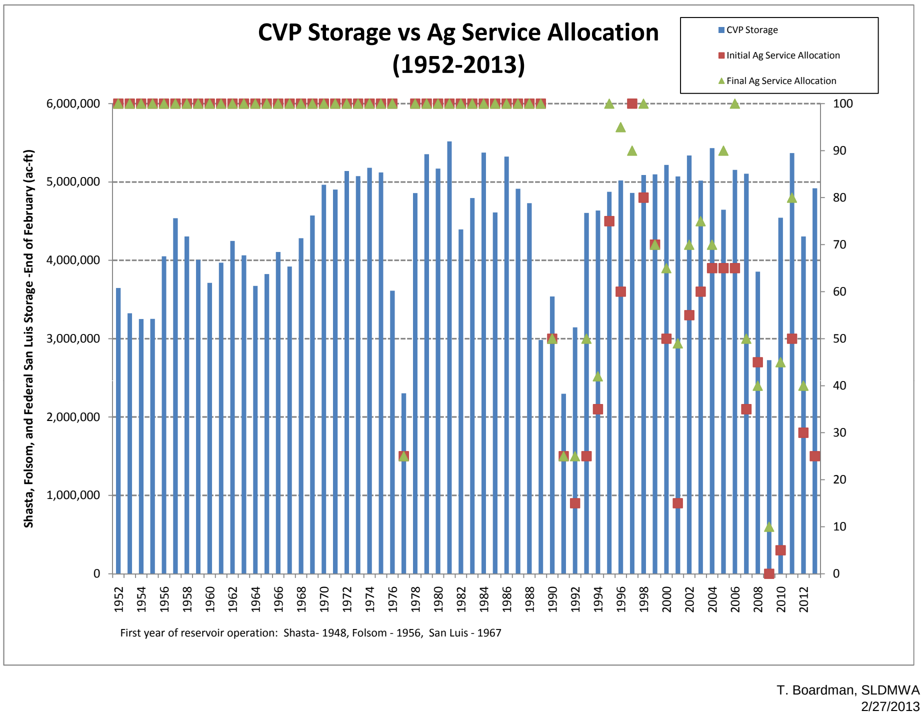

Source: Westlands irrigation district

This policy uncertainty means that valuing water rights becomes difficult. This is a problem faced by the Australian Government when purchasing rights on behalf of the environment. The future yield of these rights is highly uncertain and depends greatly on the policies of state government agencies (see Crase et al. 2009).

The second problem is determining the optimal mix of priority rights. The optimal mix will vary with changes in the composition of demand or the inflow distribution. In practice, this mix is generally fixed based on historical demands.

Freebairn and Quiggin (2006) argue this mix could be determined by market forces: via some priority right conversion market. In appendix D we show that these markets do not work in practice — because changes in the mix of rights can not occur without substantial externalities. Formal conversion of priority rights is rarely observed in practice and the few examples have been failures[3].

Another issue is enforcement. During droughts there is a temptation to allocate water on the basis of political priorities rather than official priorities. This can be a problem when urban and agricultural water demands are in conflict. For example, in the western US, low priority agricultural rights purchased by cities, often become effective high priority, given the political rule that cities come before farmers during shortages[4].

We also need to acknowledge the problem of path dependence. Once entrenched, priority rights are difficult to remove, because any changes in priority lead to large welfare effects. Further, these effects can be hard to value and therefore compensate for.

6.5 Solving the model

For this chapter we introduce some minor changes to the computational methods employed in chapter 5.

6.5.1 Inflow shares

Central case

The release sharing scenarios are solved (for the social welfare maximising storage policy) using single agent fitted Q-V iteration (see chapter 8).

To estimate the optimal inflow shares we use a simple local search method: stochastic hill climbing (see appendix B). In the release sharing case, the model is solved repeatedly for varying values of Λ_{high}. In the capacity sharing scenarios, we adjust the inflow shares in-between iterations of the learning algorithm. Figures 6.6 and 6.7 show the results of inflow share searches for two release sharing scenarios.

[3] Conversion of general and high security water rights in the Murrumbidgee region was abolished in 2008 after it led to large externalities (NSWIC 2014).

[4] In the short term allocating water to cities before farmers may well be efficient, given absent / weak spot markets. However, in the long term it creates policy uncertainty over the true yield / value of each water right class.

General case

For the sensitivity analysis we adopt the same approach to inflow shares as chapter 5.

In appendix D we present the results of 200 model runs for each of the CS-O, RS-O, CS-HL-O and RS-HL-O scenarios. From the results of these runs we fit a regression model, which predicts the optimal Λ_{high} values conditional on the parameters.

Given estimates of $\hat{\Lambda}_{high}$ we define Λ_{high} as

$$\Lambda_{high} = N_0^1(\hat{\Lambda}_{high}, 0.025)$$

6.5.2 Risk aversion

The other addition to the model is risk aversion. Here we add a 'negative exponential' utility function v_h over user payoffs u_{it}

$$v_h(u_{it}) = \begin{cases} 1 - e^{-\psi u_{it}} & \text{if } \psi \neq 0 \\ u_{it} & \text{otherwise} \end{cases}$$

$$u_{it} = \pi_h(q_{it}, \tilde{I}_t, e_{it}) + P_t(a_{it} - q_{it})$$

The negative exponential form was chosen because it can handle negative profit levels. Risk preferences are parametrised as follows

$$\bar{\pi}\psi \sim \begin{cases} 0 & \text{with probability } 0.5 \\ U[0,3] & \text{with probability } 0.5 \end{cases}$$

$$\bar{\pi} = \frac{\pi_{low}(\bar{q}_{it}, 1, 1) + \pi_{high}(\bar{q}_{it}, 1, 1)}{2}$$

That is, we assume relative risk aversion (for an average farm, in a mean inflow year) of between 0 and 3. However, we take 0 (risk neutrality) as our central case assumption. Figure 6.2 shows the shape of the users' utility functions v_h for varying levels of risk aversion (with utility scaled to a maximum of 1). Figure 6.3 shows the effect of risk aversion on the shape of the social welfare function.

Note that by introducing risk aversion into the model, we are implicitly assuming some imperfection in financial markets: that is users can not mitigate their income volatility through borrowing, insurance or other financial instruments.

Figure 6.2: User utility functions $\mathcal{R}_h$ for relative risk aversion of 0, 1.5 and 3

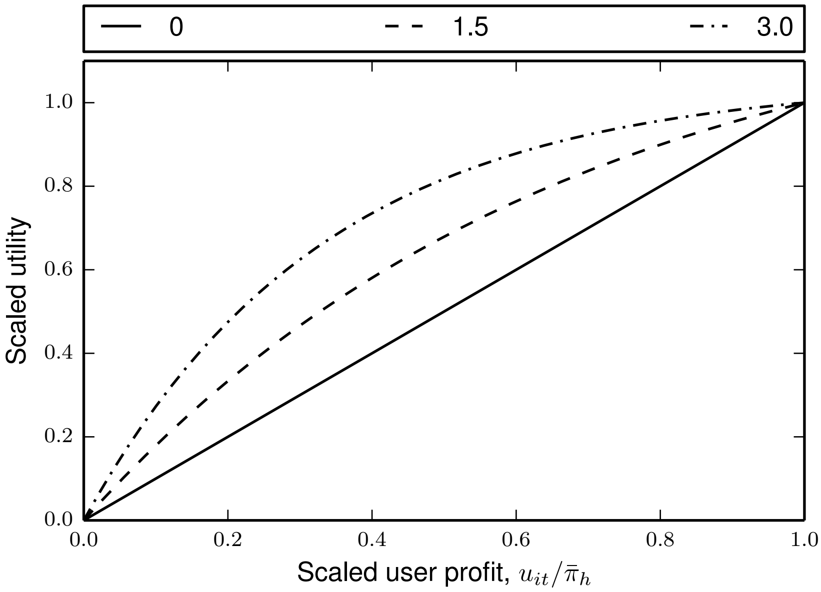

Figure 6.3: Social welfare function for relative risk aversion of 0, 1.5 and 3

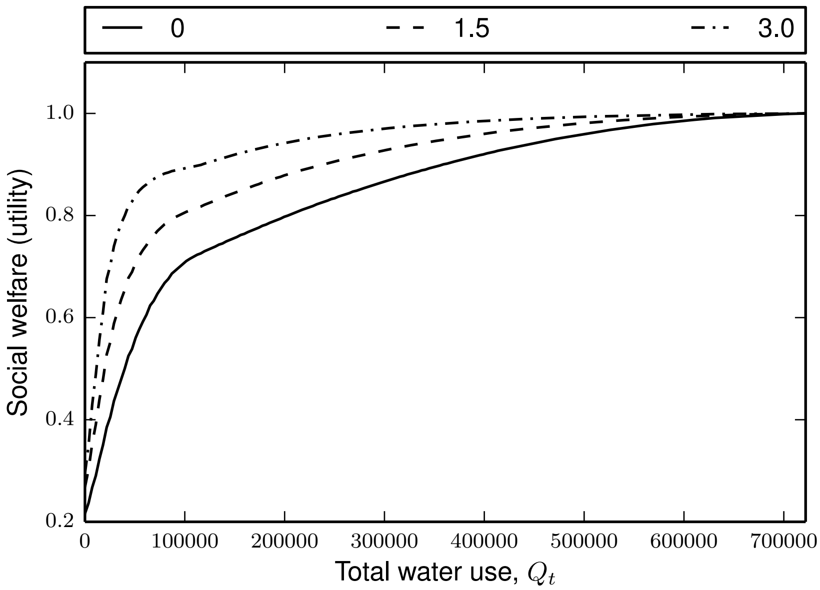

6.6 Results

6.6.1 No spot market

To get a basic understanding of the results we begin with a no trade case (the central case parameters with $\tau = \infty$). The welfare results are summarised in figure 6.4 and table 6.2.

Figure 6.4: Mean social welfare ($m) (no trade)

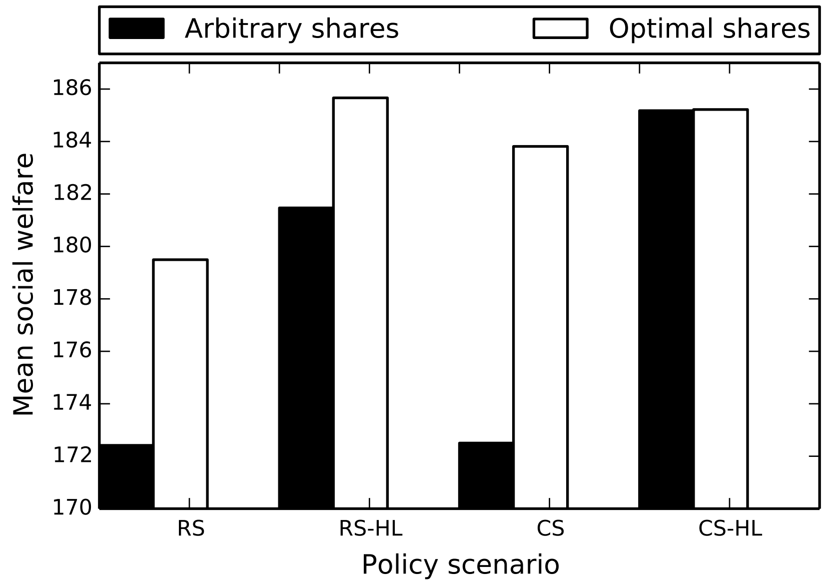

With no spot market, priority rights achieve a significant welfare gain over proportional rights. The difference in welfare between RS-HL-O and an optimal outcome — the potential gain from trade — is only around 0.5 per cent. Clearly, simple priority rights do a very good job of mimicking a post trade outcome.

With no trade, the performance of the scenarios — especially RS and CS — depends greatly on the share endowments. Figures 6.6 and 6.7 demonstrate the share search method for the release sharing scenarios. The final shares are in figure 6.5.

With proportional rights, high reliability users require large shares to avoid shortages in dry years. This means receiving excess allocation in wet years which can't be sold. With priority rights, high reliability users require much smaller shares — both under RS and CS.

Figure 6.5: High reliability inflow share Λ_{high}

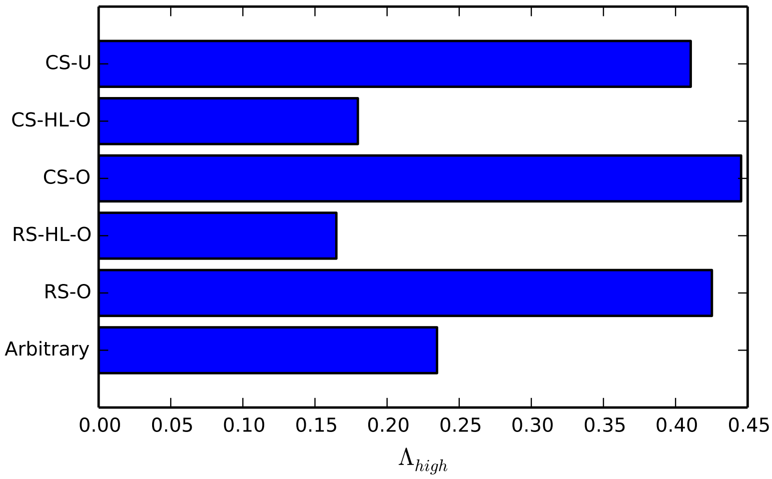

Figure 6.6: Inflow share search, RS-O

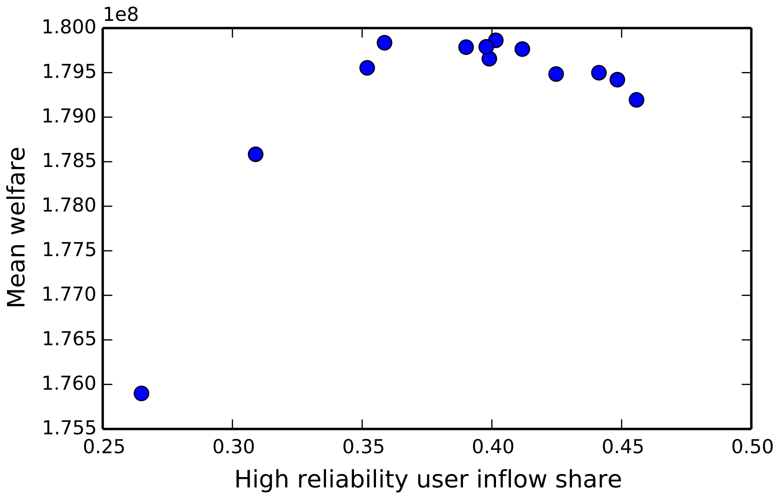

Figure 6.7: Inflow share search, RS-HL-O

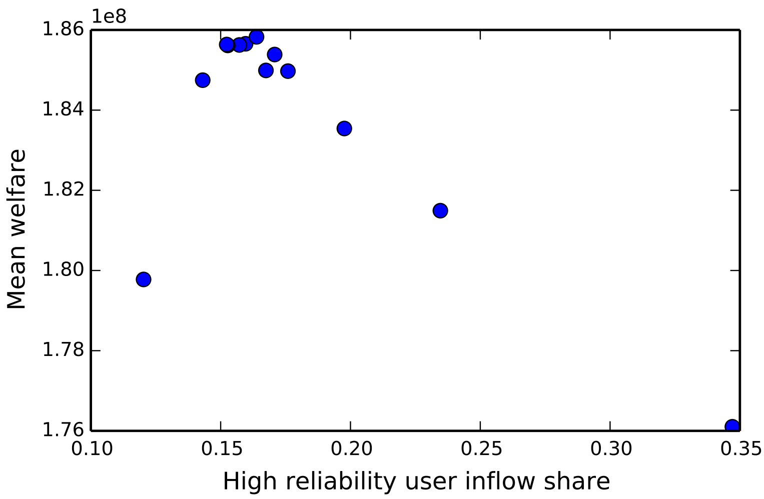

In this no trade case, RS-HL-O is the preferred scenario, followed closely by CS-HL-O. However, the performance of RS-HL-O is highly dependent on achieving the optimal level of Λ_{high} (which may be difficult in practice, see appendix D). In contrast, CS-HL performs well with or without optimal inflow shares (CS-HL outperforms RS-HL).

The results also show how storage rights mitigate trade requirements: with capacity sharing (CS-O) significantly outperforming release sharing (RS-O) in the absence of priority rights. Under capacity sharing users adopt specialised storage policies: high reliability users hold larger account balances (see tables 6.4 and 6.5).

The CS-U scenario — where inflow and capacity shares are unbundled — achieves no improvement over CS-O[5]. Under CS-U we estimate a final Λ_K of 0.42 and Λ_I of 0.42. The results suggest that equal inflow and capacity shares are essentially optimal (because unequal shares increase can storage externalities).

Finally, storage policies (both planner and user) depend on the nature of water property rights (table 6.3). With proportional rights (CS-O and RS-O) storage policy is more conservative: because large releases allocate excess water to high reliability users holding large shares. As would be expected, RS scenarios involve slightly higher storage levels than CS: because of internal spill externalities (see chapter 5).

Table 6.2: Social welfare $\sum_{i=1}^{n} u_{it}$ (\$m)

	Mean	SD	2.5th	25th	75th	97.5th
Planner	186.54	25.83	128.41	178.85	202.21	209.41
RS-HL-O	185.66	25.53	131.64	175.75	201.33	212.18
RS-HL	181.47	27.13	122.59	159.17	200.58	209.08
RS-O	179.49	28.42	80.60	180.92	192.06	198.99
RS	172.41	44.00	24.99	161.27	198.28	209.35
CS	172.50	43.67	46.93	161.05	200.40	208.56
CS-O	183.81	24.61	115.51	175.93	198.44	207.13
CS-HL	185.18	26.47	121.49	174.19	202.28	210.15
CS-HL-O	185.22	30.38	92.75	177.19	203.47	211.35
CS-U	183.63	27.77	95.77	177.72	199.44	207.38

[5] The welfare difference between CS-O and CS-U varies between model runs and occasionally CS-O is preferred. In theory, CS-U should at least match CS-O, in practice it can be lower due to sample error

Table 6.3: Storage S_t (GL)

	Mean	SD	2.5th	25th	75th	97.5th
Planner	697.08	282.63	160.05	461.25	1,000.00	1,000.00
RS-HL-O	698.61	281.44	166.85	459.28	1,000.00	1,000.00
RS-HL	671.77	286.42	152.75	425.17	1,000.00	1,000.00
RS-O	744.20	265.30	151.46	557.32	1,000.00	1,000.00
RS	692.16	291.06	120.27	453.20	1,000.00	1,000.00
CS	668.43	294.88	119.69	418.60	1,000.00	1,000.00
CS-O	733.62	261.81	186.58	530.31	1,000.00	1,000.00
CS-HL	678.30	286.95	151.57	434.04	1,000.00	1,000.00
CS-HL-O	665.79	297.48	120.58	410.11	1,000.00	1,000.00
CS-U	717.16	272.60	162.98	500.66	1,000.00	1,000.00

Table 6.4: Low reliability user storage account balance, $s_{it}/\lambda_i K$ (%)

	Mean	SD	2.5th	25th	75th	97.5th
CS	0.63	0.32	0.05	0.35	1.00	1.00
CS-O	0.62	0.33	0.06	0.34	1.00	1.00
CS-HL	0.57	0.38	0.00	0.21	1.00	1.00
CS-HL-O	0.58	0.37	0.00	0.23	1.00	1.00
CS-U	0.61	0.34	0.05	0.31	1.00	1.00

Table 6.5: High reliability user storage account balance, $s_{it}/\lambda_i K$ (%)

	Mean	SD	2.5th	25th	75th	97.5th
CS	0.68	0.31	0.05	0.42	1.00	1.00
CS-O	0.81	0.25	0.15	0.69	1.00	1.00
CS-HL	0.94	0.16	0.39	0.99	1.00	1.00
CS-HL-O	0.94	0.16	0.36	0.98	1.00	1.00
CS-U	0.74	0.21	0.15	0.65	0.88	0.88

6.6.2 Central case

The central case welfare results are summarised in figure 6.8 and table 6.6.

Figure 6.8: Mean social welfare ($m)

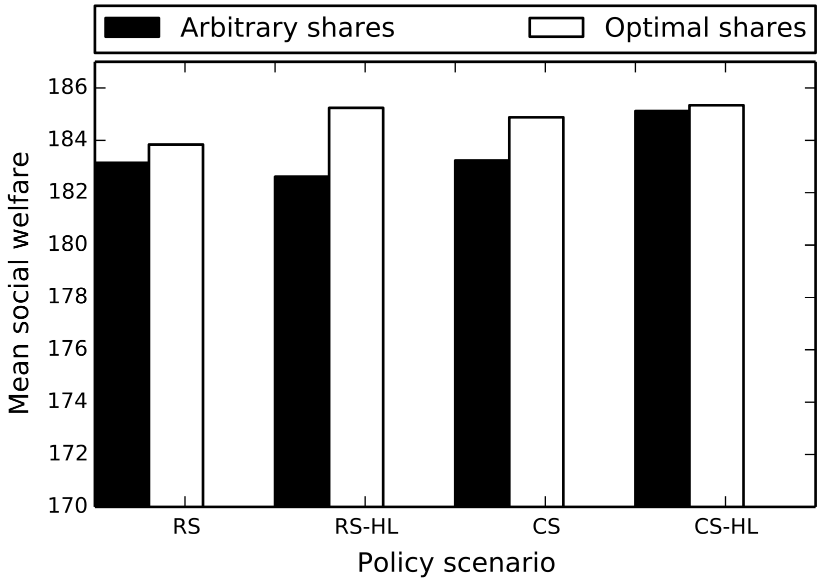

In the presence of a spot market (subject to a transfer cost) the gains from priority rights are much smaller.

Again RS-HL-O is the best scenario on mean welfare grounds. However, RS-HL is now the worst performing scenario. This is concerning given the practical constraints on achieving the optimal mix of priority rights. Once again CS-HL-O performs similarly with and without optional inflow shares.

In the central case, capacity sharing without priority rights CS-O performs almost as well as CS-HL-O. While the combination of storage and priority rights (CS-HL, CS-HL-O) greatly reduces trade requirements (table 6.8), it also induces below optimal storage volumes (table 6.7). Priority inflow rights mean that high reliability accounts fill frequently (table 6.12) exacerbating internal spill externalities. So with capacity sharing, priority rights seem to improve use allocation at the expense of inter-temporal efficiency.

Finally, scenarios with priority rights tend to favour the high reliability users at the expense of low (tables 6.9 and 6.10).

Figure 6.9: High reliability inflow share Λ_{high}

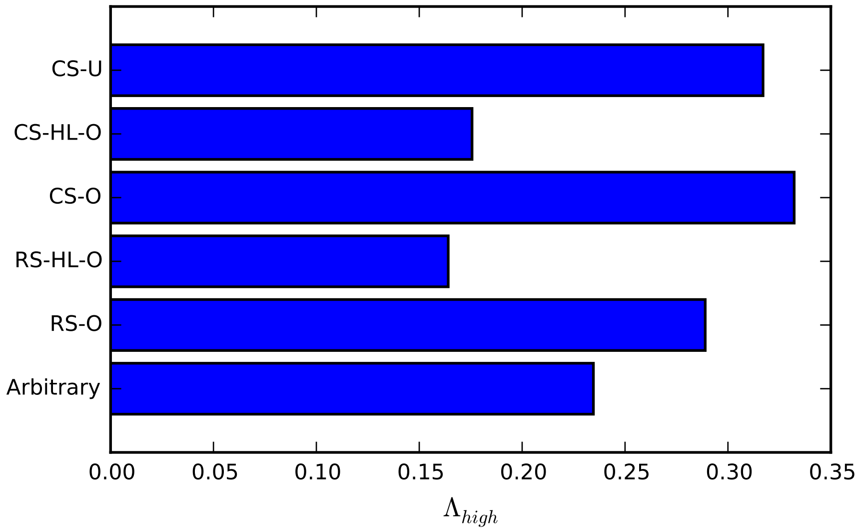

Table 6.6: Social welfare $\sum_{i=1}^{n} u_{it}$ (\$m)

	Mean	SD	2.5th	25th	75th	97.5th
Planner	186.69	25.66	129.32	179.14	202.25	209.45
RS-HL-O	185.51	23.88	143.00	176.04	201.36	211.89
RS-HL	182.61	24.57	138.82	170.39	198.96	208.20
RS-O	183.84	27.34	114.10	173.42	200.57	206.95
RS	183.14	29.80	99.65	172.09	202.31	210.19
CS	183.23	26.88	115.64	173.36	200.51	208.01
CS-O	184.88	25.85	119.55	176.88	200.86	208.66
CS-HL	185.12	26.33	119.33	174.91	201.84	209.44
CS-HL-O	185.34	28.52	103.69	176.93	202.68	210.63
CS-U	184.75	25.88	121.27	175.50	201.29	208.92

Table 6.7: Storage S_t (GL)

	Mean	SD	2.5th	25th	75th	97.5th
Planner	698.71	282.12	161.56	463.21	1,000.00	1,000.00
RS-HL-O	701.85	279.15	157.41	471.09	1,000.00	1,000.00
RS-HL	694.01	276.99	178.60	459.54	1,000.00	1,000.00
RS-O	704.33	284.00	147.75	471.65	1,000.00	1,000.00
RS	680.97	290.70	129.47	439.22	1,000.00	1,000.00
CS	695.60	281.46	156.29	462.46	1,000.00	1,000.00
CS-O	699.54	279.60	157.57	469.97	1,000.00	1,000.00
CS-HL	679.71	288.19	148.02	433.90	1,000.00	1,000.00
CS-HL-O	674.01	293.57	130.18	423.88	1,000.00	1,000.00
CS-U	699.42	279.28	161.12	469.34	1,000.00	1,000.00

Table 6.8: Absolute trade volume, $\sum_{i=1}^{n} |a_{it} - q_{it}|$ (GL)

	Mean	SD	2.5th	25th	75th	97.5th
Planner	0.00	0.00	0.00	0.00	0.00	0.00
RS-HL-O	8.28	7.83	0.00	1.31	15.21	24.16
RS-HL	68.71	42.82	0.00	16.60	102.14	109.14
RS-O	41.64	45.35	0.00	7.30	73.96	139.16
RS	50.49	52.29	0.00	4.86	95.66	154.67
CS	58.25	51.29	6.42	13.61	98.90	161.05
CS-O	26.50	33.37	2.31	6.30	30.88	120.28
CS-HL	24.64	12.50	1.46	13.71	34.95	40.96
CS-HL-O	11.55	10.22	1.00	3.94	17.11	37.05
CS-U	27.85	35.74	1.35	5.11	35.36	127.74

Table 6.9: Low reliability user welfare, $\sum \mathcal{U}_{low}$ (\$m)

	Mean	SD	2.5th	25th	75th	97.5th
Planner	79.29	18.54	40.46	68.30	92.82	96.91
RS-HL-O	77.02	20.00	39.89	64.84	92.12	99.23
RS-HL	73.69	20.58	39.78	55.95	90.90	95.93
RS-O	82.18	11.93	57.04	73.49	90.66	95.09
RS	84.30	12.29	57.88	75.93	93.68	98.16
CS	84.07	12.07	58.12	75.89	93.23	97.31
CS-O	81.30	13.97	52.21	70.82	92.60	97.07
CS-HL	77.58	19.69	40.04	62.67	92.68	96.90
CS-HL-O	79.94	19.03	40.26	68.14	93.84	98.37
CS-U	81.43	13.89	52.41	70.80	92.65	97.03

Table 6.10: High reliability user welfare, $\sum \mathcal{U}_{high}$ (\$m)

	Mean	SD	2.5th	25th	75th	97.5th
Planner	107.39	13.55	88.59	106.74	112.18	115.99
RS-HL-O	108.48	10.70	96.74	107.33	112.60	116.33
RS-HL	108.92	12.24	95.69	107.74	113.65	117.88
RS-O	101.66	19.10	34.13	99.57	111.03	115.55
RS	98.84	21.34	17.77	97.99	109.29	114.18
CS	99.16	17.67	38.82	97.55	108.35	113.24
CS-O	103.59	16.43	50.18	101.84	111.06	115.40
CS-HL	107.54	13.38	78.48	106.83	112.76	116.55
CS-HL-O	105.40	14.89	62.95	105.16	111.02	114.99
CS-U	103.33	16.43	49.92	101.57	110.90	115.21

Table 6.11: Low reliability user account balance, $s_{it}/\lambda_i K$ (%)

	Mean	SD	2.5th	25th	75th	97.5th
CS	0.64	0.32	0.06	0.37	1.00	1.00
CS-O	0.63	0.32	0.05	0.35	1.00	1.00
CS-HL	0.56	0.38	0.00	0.20	1.00	1.00
CS-HL-O	0.59	0.36	0.00	0.26	1.00	1.00
CS-U	0.63	0.32	0.05	0.35	0.99	0.99

Table 6.12: High reliability user account balance, $s_{it}/\lambda_i K$ (%)

	Mean	SD	2.5th	25th	75th	97.5th
CS	0.71	0.28	0.16	0.48	1.00	1.00
CS-O	0.73	0.29	0.12	0.51	1.00	1.00
CS-HL	0.93	0.17	0.31	0.99	1.00	1.00
CS-HL-O	0.93	0.15	0.40	0.95	1.00	1.00
CS-U	0.72	0.28	0.12	0.50	1.00	1.00

Risk aversion

Here we assume central case parameters, except for $\bar{\pi}_h \psi_h = 3$. The welfare results are shown in 'utility' units (figure 6.10, table 6.13).

Figure 6.10: Mean social welfare (utility) (risk aversion)

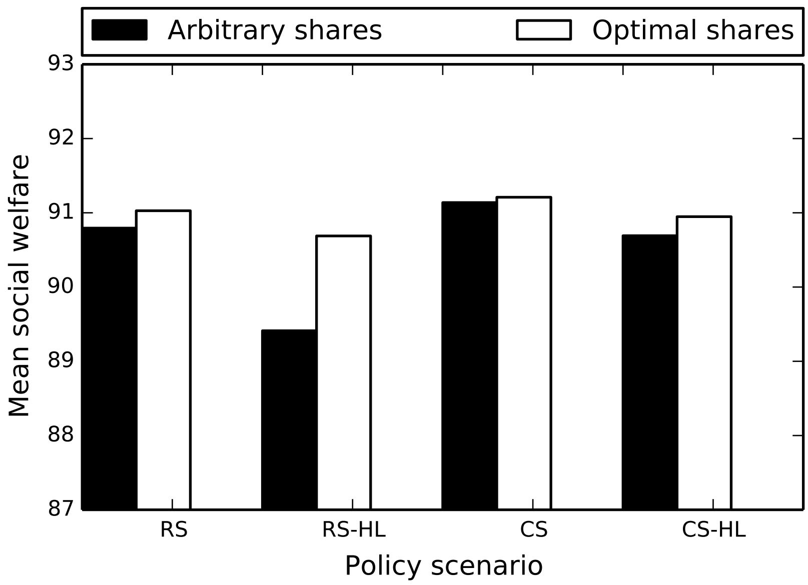

With risk aversion, it is possible for decentralised scenarios to achieve higher utility than the planner's solution. While the planner's allocation of water is optimal, risk may not be optimally shared between users — there may be a welfare gain from some form of income redistribution.

Figure 6.11: High reliability inflow share Λ_{high} (risk aversion)

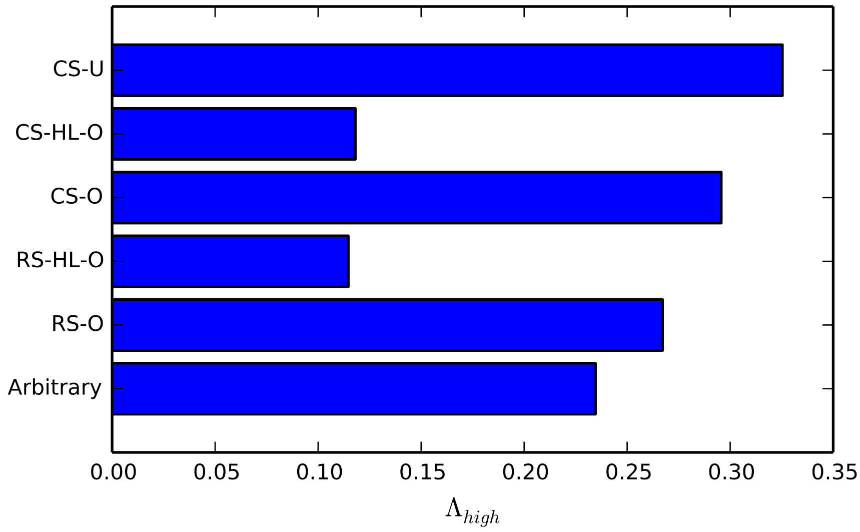

In the decentralised case, users have a second income stream — water trading — which can help to redistribute income risk. This second income stream is of particular benefit to low reliability users, as it allows them to maintain their income in dry years by selling allocations to high reliability users. As a result of these effects, low reliability users demand larger shares in these scenarios (figure 6.11).

With risk aversion the preferred scenario is CS-O and priority rights are generally outperformed by proportional rights[6]. From a risk management perspective, priority rights worsen the position of low reliability users (table 6.16). So while priority rights may convey a risk benefit to their holders, they impose a risk cost on the low priority users, leading to little if any net gain.

Finally, as would be expected, optimal storage policy is more conservative across all scenarios when users are risk averse (table 6.14).

[6] This result is dependent on our restriction that only high reliability users can hold priority rights.

Table 6.13: Social welfare $\sum_{i=1}^{n} u_{it}$ (utility)

	Mean	SD	2.5th	25th	75th	97.5th
Planner	90.33	5.71	79.30	89.81	92.94	93.54
RS-HL-O	90.69	4.17	83.16	90.04	92.86	93.39
RS-HL	89.41	5.19	81.18	85.42	92.81	93.67
RS-O	91.03	5.18	84.88	90.48	93.09	93.58
RS	90.79	3.82	85.34	89.13	93.02	93.59
CS	91.14	5.24	82.29	90.75	93.20	93.76
CS-O	91.21	5.11	83.65	90.63	93.32	93.88
CS-HL	90.69	6.28	80.10	90.74	93.34	93.97
CS-HL-O	90.95	6.11	79.20	91.21	92.99	93.59
CS-U	91.19	5.23	83.79	90.56	93.37	93.99

Table 6.14: Storage S_t (GL)

	Mean	SD	2.5th	25th	75th	97.5th
Planner	762.24	256.02	201.61	575.54	1,000.00	1,000.00
RS-HL-O	800.96	225.79	298.41	638.61	1,000.00	1,000.00
RS-HL	781.16	232.60	288.70	601.02	1,000.00	1,000.00
RS-O	782.07	241.73	255.82	607.30	1,000.00	1,000.00
RS	821.68	206.52	342.54	675.43	1,000.00	1,000.00
CS	750.93	257.46	205.80	556.06	1,000.00	1,000.00
CS-O	749.36	256.05	213.55	553.65	1,000.00	1,000.00
CS-HL	725.96	270.46	178.47	512.47	1,000.00	1,000.00
CS-HL-O	746.14	268.53	167.38	542.84	1,000.00	1,000.00
CS-U	745.53	257.40	212.12	546.12	1,000.00	1,000.00

Table 6.15: High reliability user welfare, $\sum \mathcal{U}_{high}$ (utility)

	Mean	SD	2.5th	25th	75th	97.5th
Planner	47.19	3.92	45.03	47.35	47.75	48.00
RS-HL-O	47.06	2.34	46.44	46.97	47.38	47.66
RS-HL	47.52	2.43	46.64	47.45	47.88	48.15
RS-O	46.67	4.19	44.57	46.81	47.53	47.85
RS	46.65	2.72	44.88	46.41	47.35	47.72
CS	46.35	4.14	39.71	46.49	47.41	47.73
CS-O	46.67	3.88	42.60	46.74	47.56	47.85
CS-HL	47.15	4.35	45.87	47.31	47.77	48.02
CS-HL-O	46.51	4.56	44.36	46.71	47.28	47.59
CS-U	46.77	3.96	43.05	46.85	47.62	47.91

Table 6.16: Low reliability user welfare, $\sum \mathcal{U}_{high}$ (utility)

	Mean	SD	2.5th	25th	75th	97.5th
Planner	43.15	3.34	34.18	42.16	45.49	46.19
RS-HL-O	43.63	3.13	36.28	42.86	45.66	46.22
RS-HL	41.89	4.50	33.89	37.46	45.30	46.20
RS-O	44.36	1.94	39.94	43.57	45.67	46.20
RS	44.14	1.93	40.27	42.74	45.73	46.33
CS	44.79	1.72	40.66	44.22	45.90	46.38
CS-O	44.54	2.00	39.72	43.79	45.92	46.42
CS-HL	43.54	3.64	34.12	43.12	45.92	46.36
CS-HL-O	44.44	2.52	34.83	44.23	45.85	46.31
CS-U	44.43	2.14	39.33	43.52	45.96	46.43

6.6.3 General case

Now we present the results of 1,500 model runs, for the CS, RS, CS-HL and RS-HL scenarios — each with approximately optimal inflow shares as detailed in section 6.5.1. For each run we record the following statistics:

- Social welfare: $\frac{1}{T}\sum_{t=1}^{T}\sum_{i=1}^{n} u_{it}$
- Low and high reliability welfare: $\frac{1}{T}\sum_{t=1}^{T}\sum_{i\in\mathcal{U}^{low}} u_{it}$, $\frac{1}{T}\sum_{t=1}^{T}\sum_{i\in\mathcal{U}^{high}} u_{it}$
- Storage: $\frac{1}{T}\sum_{t=1}^{T} S_{it}$

For each statistic we also define an index relative to the CS scenario. Summary statistics are presented at the end of this section in tables 6.19 to 6.23.

Social welfare

On the basis of mean social welfare CS is the best performing scenario overall (figure 6.12, table 6.19). CS is preferred in 612 of the model runs, RS-HL 408, CS-HL 241 and RS 211[7].

Next we regress the mean welfare indexes against the parameters (table 6.17, figure 6.13). The most important parameter (table 6.17) is the number of high reliability users n_{high}, followed by the ratio of inflow to capacity and the inflow coefficient of variation.

CS-HL performs better in cases with few high reliability users and RS-HL better in cases with more (figure 6.13b). Both priority rights scenarios (CS-HL and RS-HL) perform poorly with low mean inflow relative to capacity (figure 6.13a).

[7]21 runs were excluded due to numerical errors and in 9 cases the results were too close to call.

Figure 6.12: Social welfare index, general case

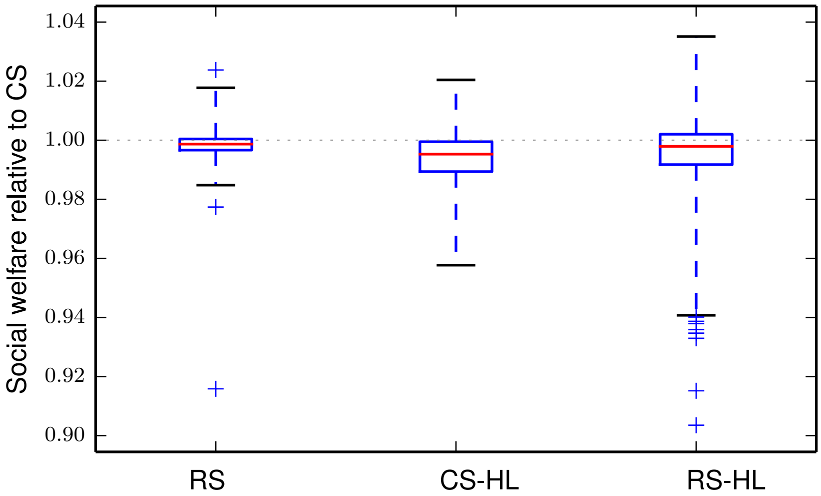

Figure 6.13: Welfare index regression results

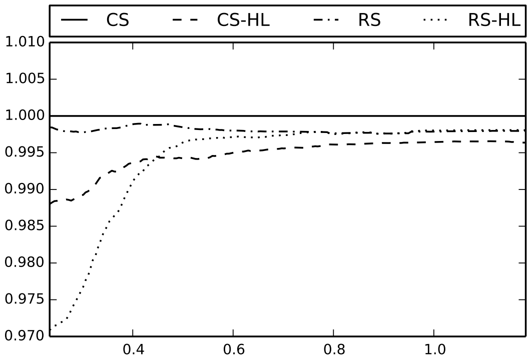

(a) Mean inflow to capacity, $E[I_t]/K$

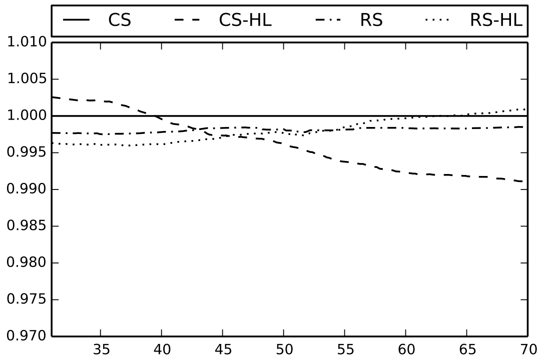

(b) Number of high reliability users, n_{high}

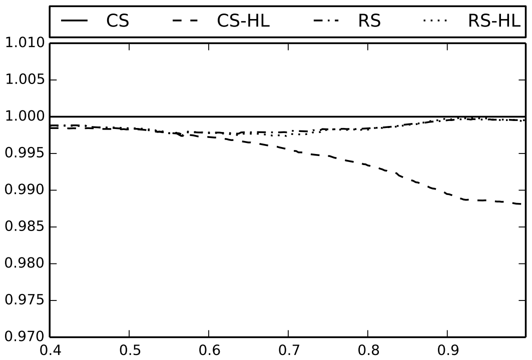

(c) Inflow variability, cv_I

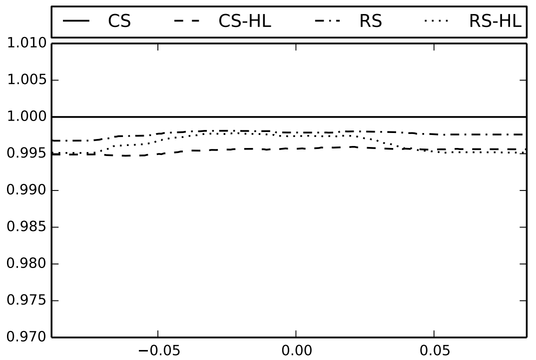

(d) RS flow share, $\Lambda^{RS}_{high} - \hat{\Lambda}^{RS}_{high}$

Next, we regress the preferred scenario (as a qualitative dependent variable) against the model parameters (table 6.18). Figures 6.14 and 6.15 plot the preferred scenario against the two most important parameters: n_{high} and $E[I_t]/K$ (once again the shaded areas represent the regression model predictions).

Here, we see that (with or without risk aversion) CS-HL is preferred in cases with few high reliability users and high mean inflow (shaded blue area). In risk averse case, ($\bar{\pi}\phi = 1$) CS is preferred for the vast majority of parameter combinations (red shaded area). RS-HL is often preferred in the risk neutral cases (green shaded area).

Figure 6.14: Preferred scenario by $E[I_t]/K$ and n_{high} (risk neutral, $\bar{\pi}\phi = 0$)

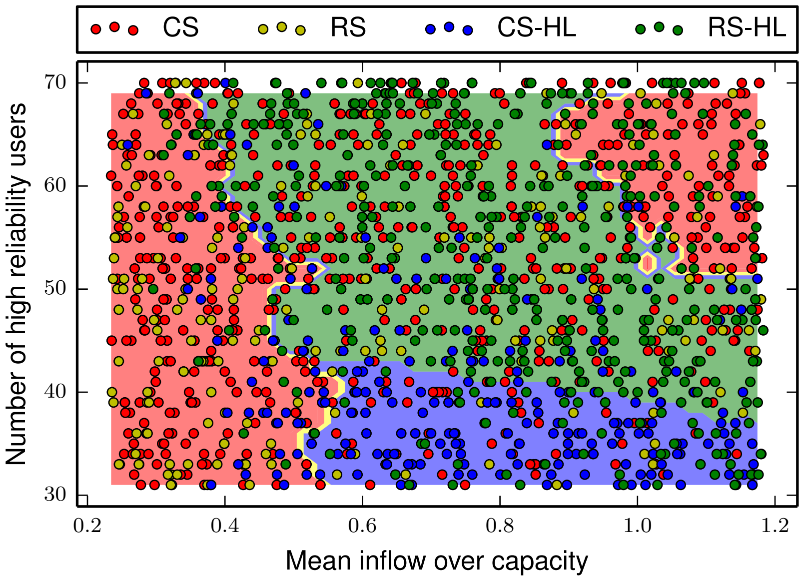

Figure 6.15: Preferred scenario by $E[I_t]/K$ and n_{high} (risk averse, $\bar{\pi}\phi = 1$)

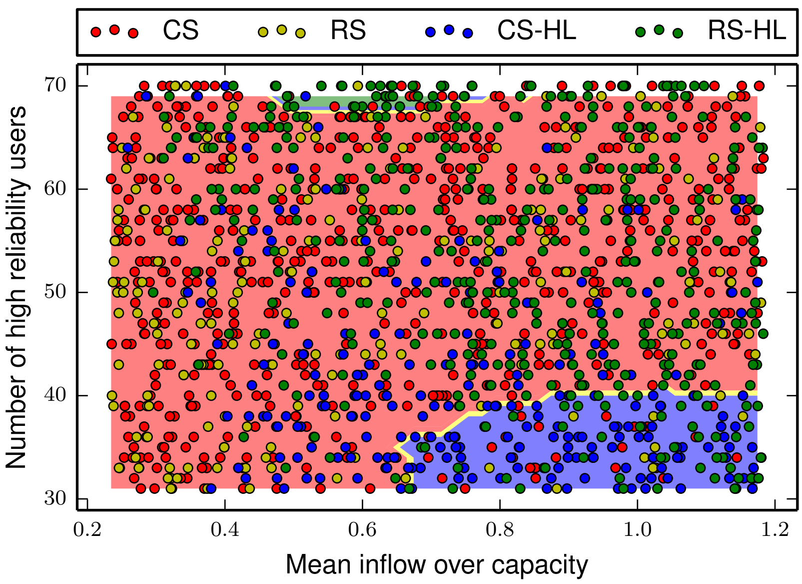

Storage

The storage results are summarised in figure 6.16, and tables 6.20 and 6.21. Here we see that RS and RS-HL generally lead to higher mean storage levels than CS, while CS-HL generally leads to lower mean storage levels — consistent with the central case results.

Figure 6.16: Storage index, general case

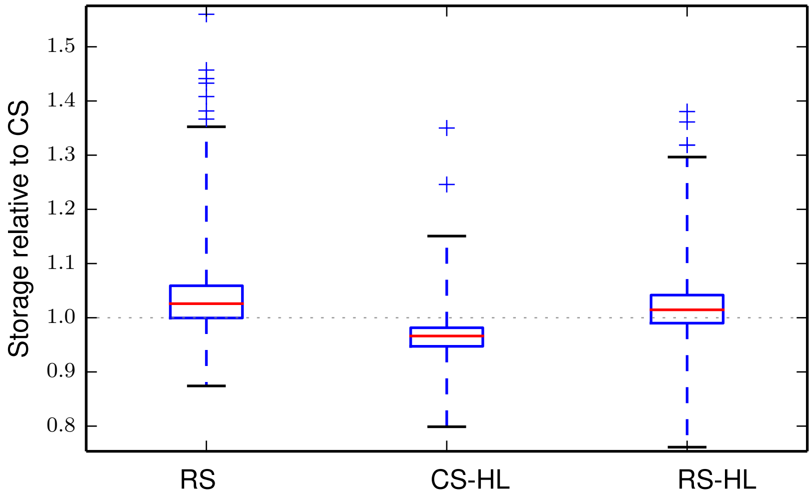

Welfare distribution

The low and high reliability user welfare results are summarised in figures 6.17 and 6.18 and tables 6.22 and 6.23. Here we see that — relative to CS — RS-HL tends to favour high reliability users over low, once again consistent with our central case results.

Figure 6.17: Low reliability user welfare, general case

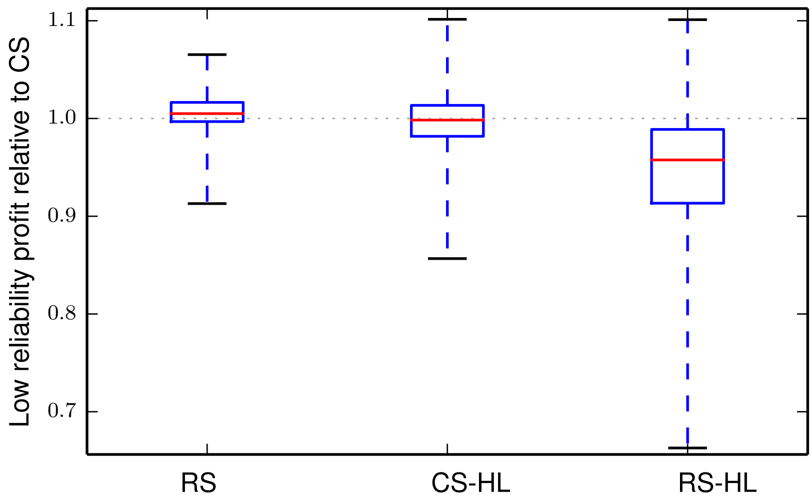

Figure 6.18: High reliability user welfare, general case

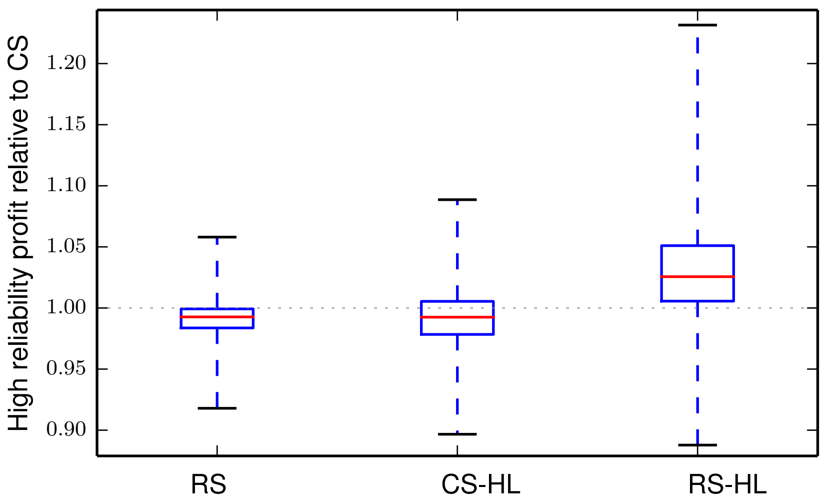

Table 6.17: Social welfare index regression, parameter importance

	Importance
$E[I]/K$	26.46
n_{high}	11.20
c_v	8.67
$\Lambda_{high}^{RS-HL} - \hat{\Lambda}_{high}^{RS-HL}$	7.47
$\Lambda_{high}^{RS} - \hat{\Lambda}_{high}^{RS}$	7.45
δ_{1a}	6.80
τ	5.91
$\bar{\pi}\phi$	4.14
$\frac{\mathscr{A}_{low}}{E[I]/K}$	3.50
$\Lambda_{high}^{C}S - \hat{\Lambda}_{high}^{CS}$	2.79
$\Lambda_{high}^{CS-HL} - \hat{\Lambda}_{high}^{CS-HL}$	2.74
ρ_I	2.39
α	2.24
σ_η	2.22
ρ_e	2.09
δ_{1b}	2.06
$\delta 0$	1.87

Table 6.18: Preferred scenario classifier, importance and sample means

	Importance	CS	RS	CS-HL	RS-HL
n_{high}	9.71	52.98	50.54	41.76	53.96
$E[I]/K$	8.92	0.66	0.62	0.75	0.81
$\bar{\pi}\phi$	7.66	0.97	0.74	0.53	0.36
c_v	7.17	0.69	0.72	0.62	0.74
δ_{1a}	6.45	49,604.85	48,746.57	45,387.31	66,482.53
$\Lambda_{high}^{RS-HL} - \hat{\Lambda}_{high}^{RS-HL}$	5.54	-0.00	0.00	0.00	0.00
$\frac{\mathscr{A}_{low}}{E[I]/K}$	5.54	6,797.63	6,802.15	6,526.14	7,006.50
$\Lambda_{high}^{RS} - \hat{\Lambda}_{high}^{RS}$	5.48	-0.00	0.00	0.00	0.00
τ	5.01	56.15	48.68	55.29	55.73
$\Lambda_{high}^{CS-HL} - \hat{\Lambda}_{high}^{CS-HL}$	5.00	0.00	-0.00	0.00	-0.00
$\Lambda_{high}^{C}S - \hat{\Lambda}_{high}^{CS}$	4.90	0.00	-0.00	0.00	-0.00
α	4.83	8.57	9.29	8.86	8.94
δ_{1b}	4.83	0.22	0.23	0.23	0.23
ρ_e	4.81	0.40	0.40	0.41	0.40
$\delta 0$	4.74	0.61	0.62	0.61	0.62
σ_η	4.72	0.15	0.15	0.15	0.15
ρ_I	4.69	0.25	0.25	0.25	0.25

Table 6.19: Social welfare index, general case

	Mean	Min	Q1	Q3	Max
CS	1.0000	1.0000	1.0000	1.0000	1.0000
RS	0.9986	0.9158	0.9967	1.0004	1.0238
CS-HL	0.9943	0.9577	0.9894	0.9995	1.0204
RS-HL	0.9952	0.9035	0.9917	1.0020	1.0351

Table 6.20: Mean storage, general case

	Mean	Min	Q1	Q3	Max
CS	691.64	326.57	629.90	768.83	898.38
RS	714.90	383.40	662.28	779.40	934.18
CS-HL	667.71	310.12	596.99	753.04	900.11
RS-HL	703.92	309.45	643.79	776.56	939.58
Planner	697.63	342.71	640.71	769.62	898.54

Table 6.21: Storage index, general case

	Mean	Min	Q1	Q3	Max
CS	1.000	1.000	1.000	1.000	1.000
RS	1.040	0.874	1.000	1.059	1.560
CS-HL	0.963	0.799	0.947	0.982	1.350
RS-HL	1.020	0.761	0.990	1.042	1.380

Table 6.22: Low reliability user welfare index, general case

	Mean	Min	Q1	Q3	Max
CS	1.000	1.000	1.000	1.000	1.000
RS	1.007	0.913	0.997	1.016	1.065
CS-HL	0.996	0.857	0.982	1.013	1.101
RS-HL	0.947	0.663	0.913	0.989	1.101

Table 6.23: High reliability user welfare index, general case

	Mean	Min	Q1	Q3	Max
CS	1.000	1.000	1.000	1.000	1.000
RS	0.991	0.918	0.984	0.999	1.058
CS-HL	0.991	0.897	0.978	1.005	1.089
RS-HL	1.032	0.888	1.006	1.051	1.232

6.7 Conclusions

In the introduction we asked: are priority rights still necessary if we have well defined storage rights? Based on the above results, the short answer is: no.

It is important to remember that our release sharing scenarios contain some very optimistic assumptions. RS and RS-HL both assume that a central planner can make optimal storage decisions, while RS-HL also assumes the planner can determine the optimal mix of priority rights (in appendix D we showed that this mix of rights can not be determined by market forces). In this sense, our RS and RS-HL scenarios represent upper bound estimates[8]. Despite these assumptions, capacity sharing still outperforms RS-HL and RS more often than not and when RS-HL does outperform CS the welfare differences are small.

The results largely support the original capacity sharing model as proposed by Dudley and Musgrave (1988). Two potential criticisms of capacity sharing (as proposed by Dudley and Musgrave (1988) and implemented at St George in Queensland) are the absence of priority rights and the bundling of inflow and capacity shares. The above results show that neither criticism is warranted.

This outcome is a good example of the second best nature of regulated rivers. While intuition suggests unbundling inflow and capacity shares and / or combining CS with priority rights may be useful reforms, both tend to exacerbate storage externalities (i.e., internal spills). An exception are rivers with little high reliability demand and high inflow relative to capacity, where CS-HL performs well: because physical storage spills are more of an issue than internal spills in that environment.

The results suggest that in the absence of path dependencies (e.g., in a new irrigation system) and in the presence of a reasonably efficient spot market, standard capacity sharing rights are hard to beat. In this sense, priority rights are best seen as a hangover from a time when water trading was not possible and storage rights were unheard of.

The problem is that in most established systems priority rights are now well entrenched. Given that the welfare gains from removing priority rights (as evidenced by the differences between CS and CS-HL) are small, they may be outweighed by transition costs. In practice then, some combination of storage rights and priority rights may be unavoidable.

[8]Our assumption that users can not hold both types of entitlements may lead to a slight underestimation of the performance of RS-HL. However, on balance RS-HL is still likely to overestimate actual performance.

Chapter 7

Water rights with in-stream demands: environmental flows

7.1 Introduction

Efforts to secure environmental flows are occurring in many of the worlds heavily regulated rivers. Frequently, this requires governments to participate in water markets. 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 acquire 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.

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

Market power is another concern since EWHs — such as the CEWH — can be much larger than irrigators. Large right holders are more likely to take into account the effect they have on aggregate variables such as prices, storage volumes and spills.

In this chapter, we add a large EWH to our decentralised model of a regulated river. Our EWH is treated identically to consumptive users under the water rights framework — they hold water rights, make storage decisions and engage in the spot market — only the EWHs withdrawals remain in the river as environmental flows, rather than being extracted.

This chapter adds to the small body of literature which models the behaviour of EWHs in water markets, in particular their water trade and storage decisions (Grafton et al. 2011*a*, Heaney et al. 2011, Beare et al. 2006, Kirby et al. 2006).

However, the main focus is on the design of water rights. This chapter asks: which form of water property rights is ideal in the presence of a large EWH? That is, we consider if and how an EWH changes our conclusions in regards to storage rights (chapter 5) and flow rights (chapter 6).

This chapter proceeds as follows. First, we present a brief review of the economic literature attempting to model EWHs. Then, we detail the model and all of the changes required to accommodate environmental flows. Finally, we present the results: first for the planner case then for the decentralised version.

7.2 Literature

Previous economic studies (for example, Dudley et al. 1998, Beare et al. 2006, Kirby et al. 2006, Heaney et al. 2011, Grafton et al. 2011*a*) have focused more on the role of EWHs in acquiring, using and trading water, than on externalities and property rights. 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 combines farm level models with a river flow model and a set of environmental objectives[2]. A major focus is estimating trade-off curves, which compare environmental and irrigation benefits for different environmental endowments.

Beare et al. (2006) consider an EWH concerned with achieving 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) targets. 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).

The potential gains from short-term environmental water trade have also been considered by Kirby et al. (2006) and more recently by Ancev (2015). Both studies confirm a 'counter-cyclical' trade pattern with the EWH selling water to farmers during droughts and buying water back in wet periods.

Grafton et al. (2011*a*) present a social planner SDP model of the Murray River. Grafton et al. (2011*a*) 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*a*) present stylised optimal EWH release rules and estimate the welfare gains from optimal versus observed historical environmental flows for the Murray.

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

Truong (2011) considers the effect on consumptive users of a reduction in the storage capacity within a theoretical model. While the paper is concerned with the effects of an EWH, the model does not include any representation of environmental objectives.

Griffin and Hsu (1993) consider in-stream demands in unregulated rivers in the presence of return flows. Griffin and Hsu (1993) demonstrate that an optimal allocation can be achieved, if perfect property rights to return flows are defined and effective 'In-stream Water Districts' (i.e. local environmental managers) participate in the market.

Finally, there are a small number of experimental studies which consider EWHs, for a review see Tisdell (2010).

7.3 The model

In this chapter we return to the general form of the model from chapter 3.

Figure 7.1: An abstract regulated river system

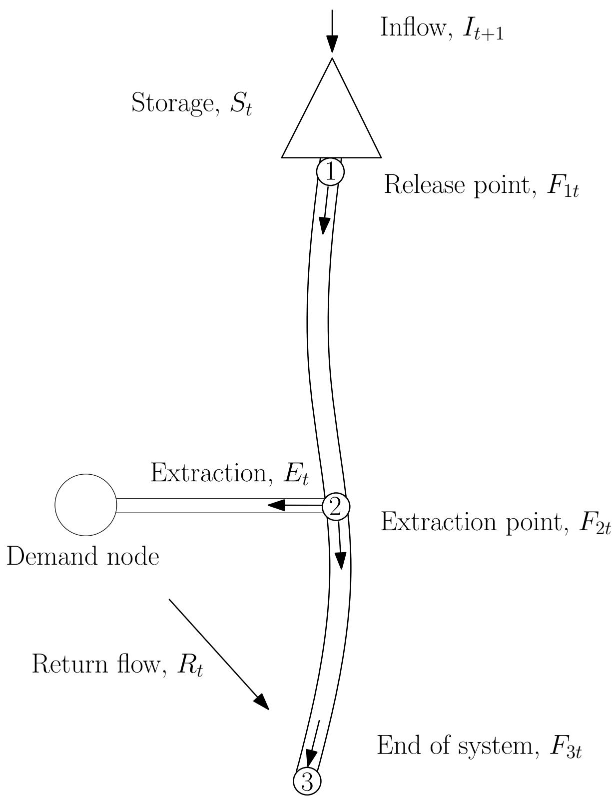

7.3.1 Inflow

This version of 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.

Seasonal inflows I_t are conditioned on an unobservable climate state variable C_t. C_t follows an *annual* AR-1 process with gamma shocks: C_t is our annual inflow model from chapter 3

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

$$0 < \rho_C < 1$$

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

Seasonal inflows 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 $[a, b]$.

7.3.2 Storage

Our storage transition rule is unchanged

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

$$0 \leq W_t \leq S_t$$

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

7.3.3 River flow

We adopt the river flow setup of our general model

$$F_{1t} = W_t + Z_t$$
$$F_{2t} = F_{1t} - \mathcal{L}_1(F_{1t}) - E_t$$
$$F_{3t} = F_{2t} - \mathcal{L}_2(F_{2t}) + R(E_t)$$

River losses are assumed fixed

$$\mathcal{L}_1(F_{1t}) = \min\{F_{1t}, \delta_{at}\}$$
$$\mathcal{L}_2(F_{2t}) = \min\{F_{2t}, \delta_{at}\}$$

Return flows are a fixed proportion of extraction

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

7.3.4 Consumptive demand

All consumptive demand (extraction) occurs in the summer period: the irrigation season. Extraction is constrained by storage withdrawals less losses

$$E_t \leq W_t - \mathcal{L}_1(W_t) = \bar{E}_t$$

As before we have $i = 1$ to N consumptive water users grouped into to low and high reliability classes. Total water use $Q_t = \sum_{i=1}^{n} q_{it}$ is constrained by extraction less delivery losses

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

Our consumptive demand model is the same as before (see chapter 3) with user profit functions $\pi_h(q_{it}, \tilde{I}_t, e_{it})$ where $\tilde{I}_t$ is defined

$$\tilde{I}_t = \frac{I_t}{E[I_t]}$$

7.3.5 Environmental demand

We adopt a simple environmental objective function, based on minimising 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 overall importance of the environment relative to consumptive users. e_{0t} reflects exogenous variation in the 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].

7.3.6 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

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

subject to the constraints detailed above.

[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).

7.4 Parameterisation

7.4.1 Inflow

The split of annual inflow between winter and summer is based on historical data for selected Southern MDB rivers (table 7.1).

Table 7.1: Winter (April-September) share of annual flow, southern MDB rivers

	Mean	SD	Min	Max
Mitta Mitta (upstream Dartmouth)	0.56	0.09	0.32	0.77
Murray (upstream Hume)	0.55	0.09	0.31	0.73
Goulburn (upstream Eildon)	0.71	0.09	0.49	0.88
Murrumbidgee (upstream Burrinjuck)	0.58	0.12	0.36	0.81

We assume $\mu_\omega \sim U[0.55, 0.75]$, $\sigma_\omega \sim U[0.09, 0.12]$ and $a_\omega = 0.30, b_\omega = 0.9$.

Storage losses

Storage losses are as before, except the evaporation rate δ_{0t} now depends on the season

$$\delta_{0t} = \begin{cases} \omega_\delta \delta_0 & \text{if } M_t = 0 \\ (1 - \omega_\delta)\delta_0 & \text{if } M_t = 1 \end{cases}$$

$$\omega_\delta \sim U[0.22, 0.36]$$

Here ω_δ reflects winter's share of annual evaporation. The distribution for ω_δ is based on average evaporation at each of our 22 MDB storages (using data from the BOM 2013).

7.4.2 Delivery losses

Our river delivery loss rates are derived largely from models of the Murray River (Gippel 2006, MDBA 2013) (see chapter 2).

For fixed delivery losses we assume

$$\frac{\delta_a}{\bar{C}} \sim U[0.02, 0.06]$$

$$\delta_{at} = \begin{cases} \omega_\delta \delta_a & \text{if } M_t = 0 \\ (1-\omega_\delta)\delta_a & \text{if } M_t = 1 \end{cases}$$

Here $\bar{C}$ is the mean annual inflow and δ_a is the annual rate of fixed loss.

Delivery losses between extraction and use are based on losses in irrigation areas in the southern MDB (see appendix A)

$$\delta_{Ea} \sim U[0, 0.1]$$
$$\delta_{Eb} \sim U[0.1, 0.3]$$

7.4.3 Return flows

We assume $\delta_R \sim U[0, 0.2]$. A return flow rate of 10 per cent is a common rule of thumb for the southern MDB (see URS Australia 2010).

7.4.4 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 that the model generates a reduction in extraction comparable with that proposed in the MDB under the Basin Plan (25.7 per cent under the planner's solution with the central case parameters, see section 7.5).

7.5 The planner's solution

We solve the planner's problem using single agent fitted Q-V iteration (chapter 8). As before, we solve the spot market with the user (and EWH) demand curves. The planner then has a problem 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 planner's problem: a 'consumptive' case where environmental benefits are ignored and an 'optimal' case where both are considered. Annual results for the central case are summarised in tables 7.3 to 7.8.

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

The optimal scenario leads to a reduction in mean extraction of 145 GL or 25.7 per cent. Mean storage and withdrawals show little change, but we do observe a significant increase in the variance of withdrawals.

Table 7.2: 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 7.3: 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 7.4: 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] As before, 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 7.5: 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 7.6: 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 7.7: 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 7.8: 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

Figure 7.2: Summer shadow water price histogram

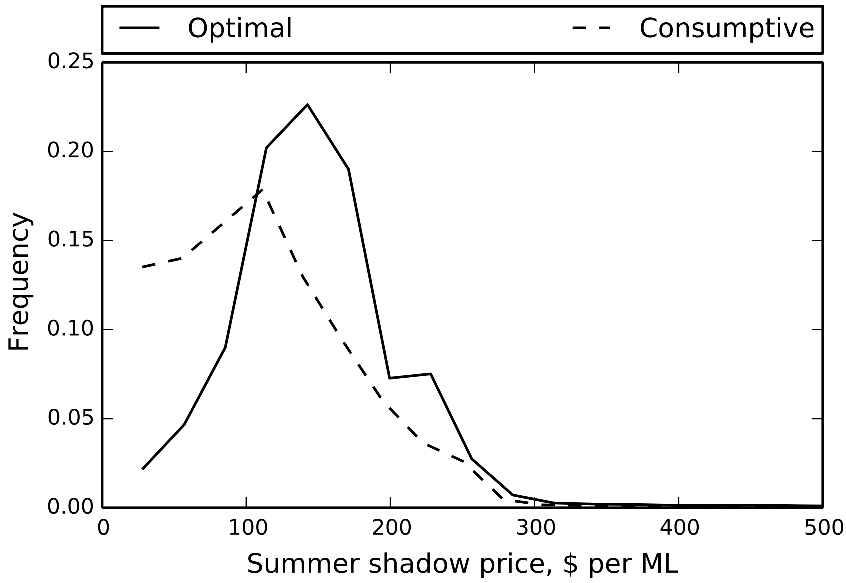

7.5.1 Environmental demands

Tables 7.9 and 7.10 detail water use and environmental flows. The majority of environmental flows are made in winter (all consumptive use occurs in summer). Environmental use is more variable than consumptive use. In many years, environmental flows are zero and in others they rival total consumptive use.

Table 7.9: Consumptive use, Q_t (GL)

	Mean	SD	2.5th	25th	75th	97.5th
Summer	296.30	122.64	84.80	198.52	395.06	504.67
Winter	0.00	0.00	0.00	0.00	0.00	0.00
Annual	296.30	122.64	84.80	198.52	395.06	504.67

Table 7.10: Environmental flow, q_{0t} (GL)

	Mean	SD	2.5th	25th	75th	97.5th
Summer	37.97	58.98	0.00	0.00	57.24	210.96
Winter	77.13	79.12	0.00	0.00	141.57	237.72
Annual	115.10	117.40	0.00	1.63	193.84	384.68

Figure 7.3 shows simulated winter storage releases (i.e., environmental flows) against winter storage and inflow. Environmental flows are mostly increasing in storage and inflow but are lower in very wet years given the occurrence of spills. The greatest demand for environmental flows occurs in high inflow but non-spill years (e.g., when a good inflow event follows a drought / low storage period) (figure 7.4).

Figure 7.3: Winter withdrawals versus storage and inflow

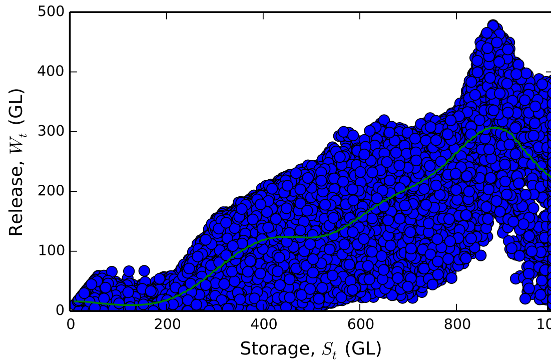

(a) Storage

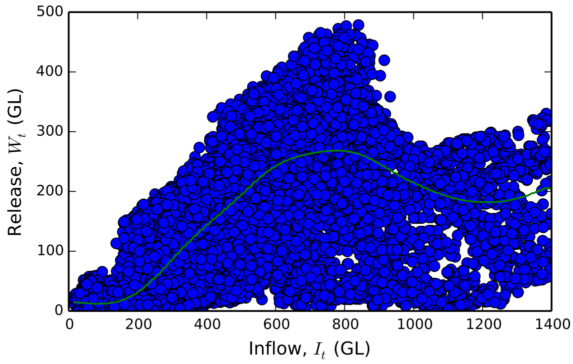

(b) Inflow

Figure 7.4: Winter withdrawals versus summer storage and winter inflow

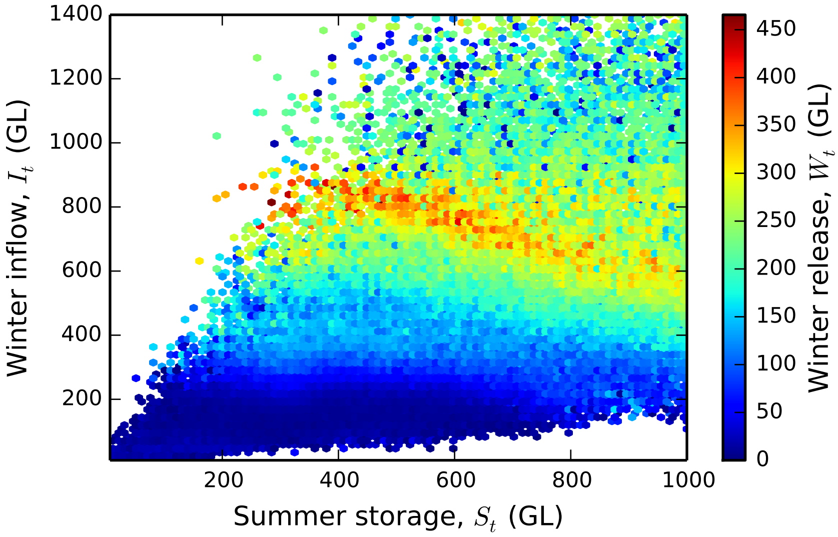

7.5.2 River flows

Next we look at total river flow, both upstream and downstream of the extraction point (tables 7.11 to 7.16) and (figures 7.5). The results show 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 relative to summer flow.

Figure 7.5 compares the natural flow distribution with those of the optimal and the consumptive scenarios — in the form of *duration curves*[6].

In winter, we see a significant increase in the frequency of small and medium flow

[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.

events, but little change in large flood events (e.g., greater than 1000 GL). This emphasis on low and medium flows is broadly consistent with environmental flows in the MDB. In addition to the high opportunity costs, the creation of very large flow events has the potential to cause flood damage in practice.

Figure 7.5: River flow duration curves, summer

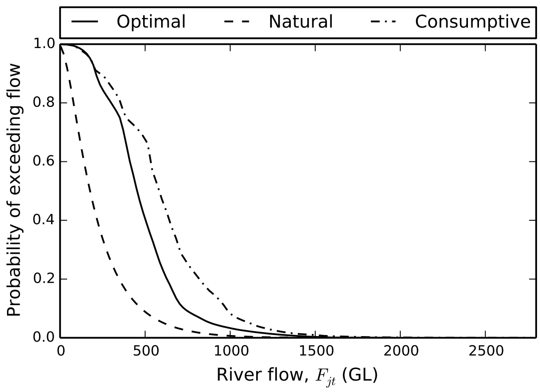

(a) Summer upstream, F_{1t}

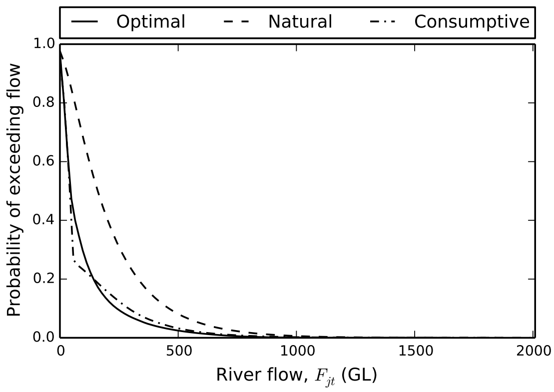

(b) Summer downstream, F_{3t}

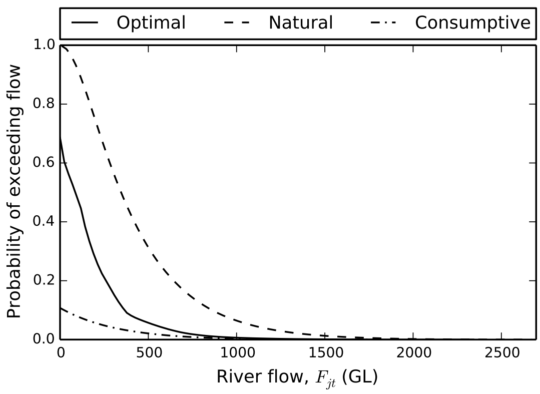

(c) Winter upstream, F_{1t}

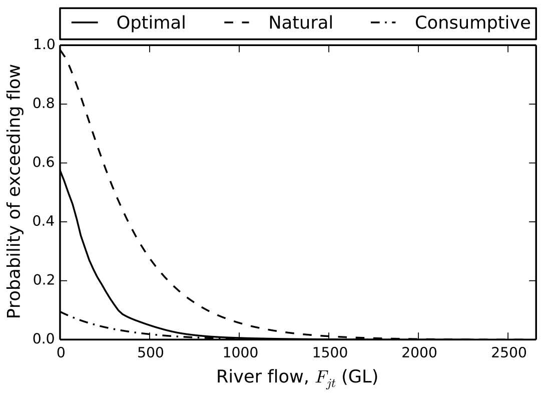

(d) Winter downstream, F_{3t}

Table 7.11: Upstream river flow, F_{1t} (GL) — natural

	Mean	SD	2.5th	25th	75th	97.5th
Summer	248.93	193.21	37.45	114.03	325.88	764.75
Winter	457.89	328.09	78.65	224.28	600.20	1,324.16
Annual	706.82	503.89	122.81	347.58	926.24	2,037.38

Table 7.12: Upstream river flow, F_{1t} (GL) — optimal

	Mean	SD	2.5th	25th	75th	97.5th
Summer	497.40	237.16	158.52	345.32	602.67	1,109.00
Winter	171.62	195.24	0.00	22.98	254.14	687.60
Annual	669.02	389.11	174.94	419.26	818.56	1,700.17

Table 7.13: Upstream river flow, F_{1t} (GL) — consumptive

	Mean	SD	2.5th	25th	75th	97.5th
Summer	631.63	285.65	165.56	415.14	770.04	1,251.02
Winter	35.59	142.87	0.00	0.00	0.00	477.88
Annual	667.22	370.00	165.56	415.14	803.25	1,641.35

Table 7.14: Downstream river flow, F_{3t} (GL) — natural

	Mean	SD	2.5th	25th	75th	97.5th
Summer	232.49	193.20	21.00	97.59	309.44	748.30
Winter	417.66	328.06	38.40	184.02	559.95	1,283.91
Annual	650.16	503.84	66.12	290.88	869.54	1,980.68

Table 7.15: Downstream river flow, F_{3t} (GL) — optimal

	Mean	SD	2.5th	25th	75th	97.5th
Summer	111.21	136.90	15.03	37.64	130.82	517.79
Winter	139.83	188.56	0.00	0.00	213.88	647.34
Annual	251.04	298.84	15.69	51.16	340.13	1,086.98

Table 7.16: Downstream river flow, F_{3t} (GL) — consumptive

	Mean	SD	2.5th	25th	75th	97.5th
Summer	113.98	148.10	15.73	40.69	76.77	558.17
Winter	31.35	134.18	0.00	0.00	0.00	437.63
Annual	145.33	245.83	15.73	40.69	111.56	881.34

7.6 The decentralised model

In the decentralised case, 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.

7.6.1 User water accounts

Our general water rights framework is the same as chapter 5, only here we allow for the possibility of priority inflow rights, as introduced in chapter 6. Each user controls their own 'water account', the evolution of account balances s_{it} follows the general form

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

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

$$\sum_{i=0}^{n} \lambda_i = 1, \quad \sum_{i=0}^{n} s_{it} = S_t, \quad \sum_{i=0}^{n} l_{it} = L_t$$

where l_{it} are user storage loss deductions, k_{it} are user account limits and x_{it} are the 'storage externalities' and g_h represents the inflow rights system: either proportional or priority.

7.6.2 Storage release rules

Aggregate withdrawals W_t depend on the sum of user withdrawals. In summer we have

$$W_t = \begin{cases} \sum_{i=0}^{N} w_{it} + 2\delta_{at} + \delta_{Ea}/(1-\delta_{Eb}) & \text{if } S_t > 2\delta_{at} + \delta_{Ea}/(1-\delta_{Eb}) \\ 0 & \text{otherwise} \end{cases}$$

Here $\delta_{at} + \delta_{Ea}/(1-\delta_{Eb})$ is the minimum release required to cover fixed losses, and the extra δ_{at} is a planned (i.e., rules based) minimum flow.

In winter W_t is just the EWH's order plus a minimum release of $2\delta_{at}$

$$W_t = \begin{cases} q_{0t}/(1-\delta_{Eb}) + 2\delta_{at} & \text{if } S_t > 2\delta_{at} + \delta_{Ea}/(1-\delta_{Eb}) \\ 0 & \text{otherwise} \end{cases}$$

For further detail on these rules see appendix B.

7.6.3 Inflow shares

In all scenarios we assume fixed inflow shares. Further we assume that rights held by the environment are low priority. As such, λ_i is defined

$$\lambda_i = \begin{cases} \Lambda^*_{high}/n_{high} & \text{if } i \in \mathcal{U}_{high} \\ \Lambda^*_{low}/n_{low} & \text{if } i \in \mathcal{U}_{low} \\ \lambda_0 & \text{if } i = 0 \end{cases}$$

where

$$\Lambda^*_{low} = \max\{1 - \Lambda_{high} - \lambda_0, 0\}$$

$$\Lambda^*_{high} = \Lambda_{high} + \min\{1 - \Lambda_{high} - \lambda_0, 0\}$$

The EWH's share λ_0 is based on the percentage change in extraction between the 'optimal' and 'consumptive' planner scenarios $\hat{\lambda}_0$ (25.7 per cent in the central case). The distribution for Λ_0 is then

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

Λ_{high} is set using our optimal share search from chapter 6. For this chapter we fit a simple linear model for $\hat{\Lambda}_{high}$ conditional on n_{high}, as shown in figure 7.6. The distributions for Λ_{high} are then defined

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

Figure 7.6: Linear model for predicted inflow share $\hat{\Lambda}_{high}$

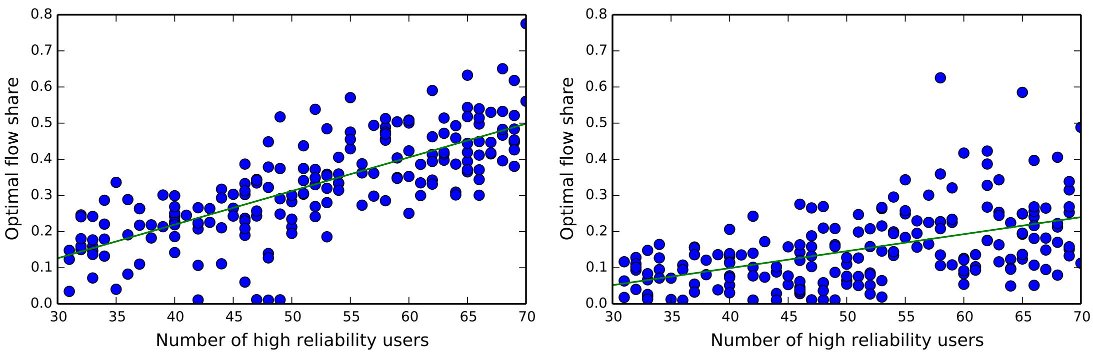

(a) Proportional right scenarios (b) Priority right scenarios

7.6.4 The spot market

The EWH is a participant in the spot market. In summer, the EWH can choose to sell allocations to users (to increase extraction) or to purchase user allocations (to reduce extraction). In winter, the EWH can purchase water allocations offered by consumptive users (to increase releases).

All rights holders receive water 'allocations' a_{it} in consumption (demand node) units[7].

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

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

$$u_{it} = \begin{cases} \pi_{ht}(q_{it}, \tilde{I}_{it}, e_{it}) + (P_t - \tau/2)(a_{0t} - q_{it}) & \text{if } a_{it} - q_{it} \geq 0 \\ \pi_{ht}(q_{it}, \tilde{I}_{it}, e_{it}) + (P_t + \tau/2)(a_{0t} - q_{it}) & \text{if } a_{it} - q_{it} < 0 \end{cases}$$

where P_t is the market price for water in consumption units and τ is the transfer cost.

Here the EWH has a payoff u_{0t}

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

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

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

Note that in winter consumptive use $\sum_{i=1}^{n} q_{it}$ and extraction E_t will both be zero. However users can still withdraw water in the hope of selling allocations to the EWH[8].

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

[8] Any of these withdrawals not purchased by the EWH — in the case where environmental flows have zero or negative marginal value — are returned to the users' accounts.

7.6.5 Users problem

As before the problem for user i 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 user's problem has one policy variable w_{it} and five state variables $s_{it}, S_t, \tilde{I}_t, e_{it}, M_t$.

7.6.6 Environmental manager's problem

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

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

The EWH's objective includes both environmental benefits and net trade proceeds / costs. As is typical in the literature we also apply a budget 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 appendix E[9].

With q_{0t} determined in the spot market, the EWH's problem, has one policy variable w_{0t} and five state variables $s_{it}, S_t, \tilde{I}_t, e_{0t}, M_t$.

7.6.7 Spot market equilibrium

In appendix E we derive the spot market equilibrium conditions. The EWH's quadratic benefit function yields demand curves which are linear in the parameters. Given these conditions the spot market can be solved as previously (see appendix B).

[9] We choose not to explicitly include the budget constraint in the EWH's problem mostly for computational reasons. Doing so 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 expect to see some 'precautionary saving' behaviour from the EWH.

7.7 Results

7.7.1 Central case

Below we present results for the policy scenarios: CS, CS-HL, SWA, SWA-HL, OA and NS, which are all as defined in chapters 5 and 6 (except for SWA-HL which is just the SWA scenario from chapter 5 with priority inflow rights).

To begin we show mean storage (figure 7.7) and mean price (figure 7.8) over the course of the learning algorithm. Similar to chapter 5 we see clear and stable differences between the policy scenarios: adding a large EWH does not damage the stability of the algorithm.

Figure 7.7: Mean storage by iteration

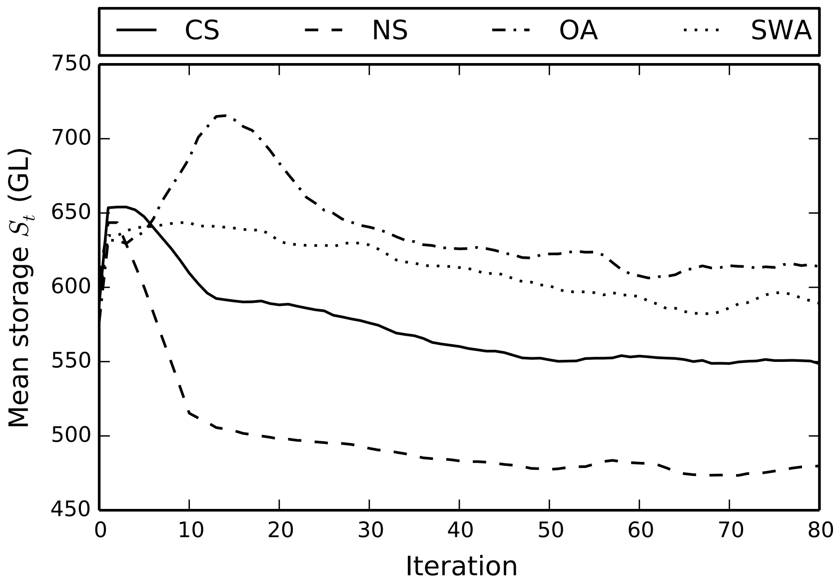

Figure 7.8: Mean price by iteration

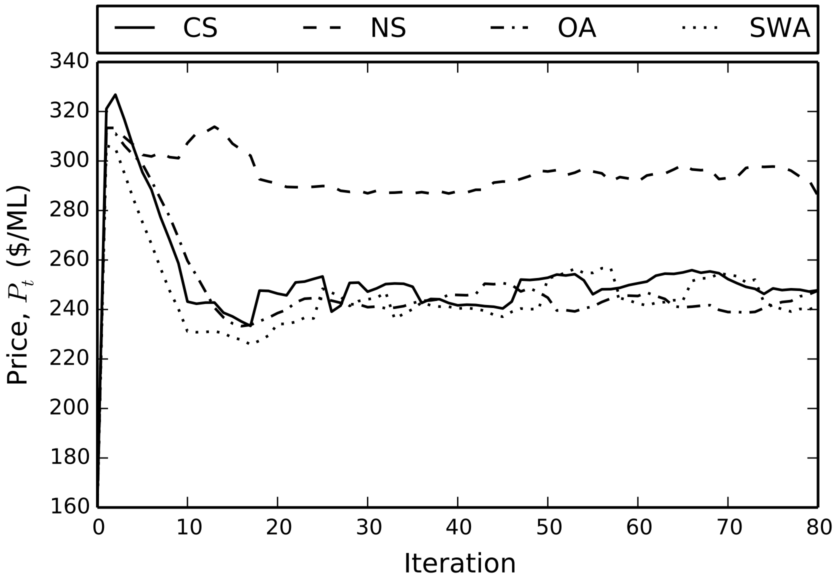

Figure 7.9: Mean environmental trade by iteration

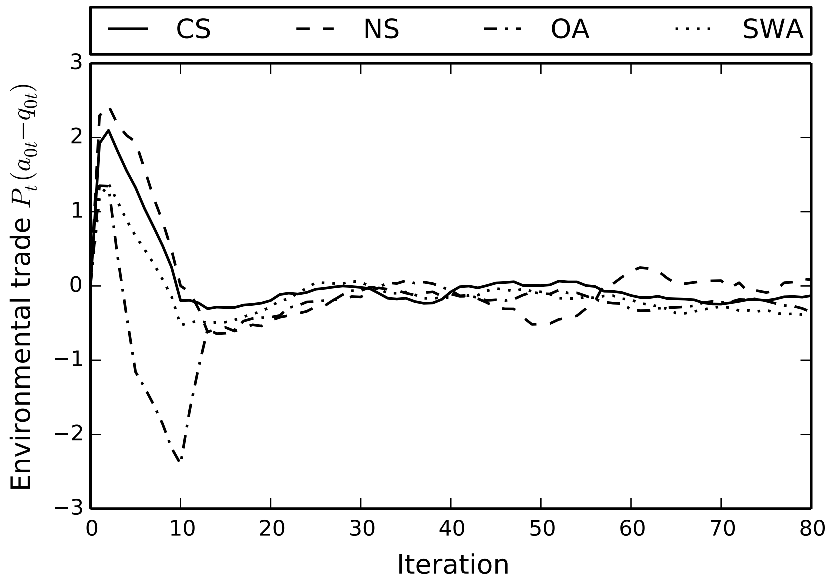

The final social welfare results are presented in table 7.17, with profits in table 7.19 and environmental benefits in table 7.20[10]. Overall, CS-HL delivers the highest social welfare ($213.8m) and NS the lowest ($209.3m). All of the decentralised scenarios remain some distance from the planner's solution at $217m.

In terms of storage levels we see a similar outcome to chapter 5 (table 7.18). OA leads to the highest storage levels, NS the lowest and once again SWA leads to 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 (table 7.18). This result is easily understood in terms of 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 — largely by accident — results in storage levels close to optimal (which explains why OA performs reasonably well in terms of social welfare).

Table 7.17: 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 7.18: 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

[10]Here we focus on annual results, some seasonal results are presented in appendix E.

Table 7.19: 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

Table 7.20: 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 7.21: 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 7.22: 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 7.23: 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 7.10: Mean profit versus mean environmental benefits, $\lambda_{0t} = 0.263$

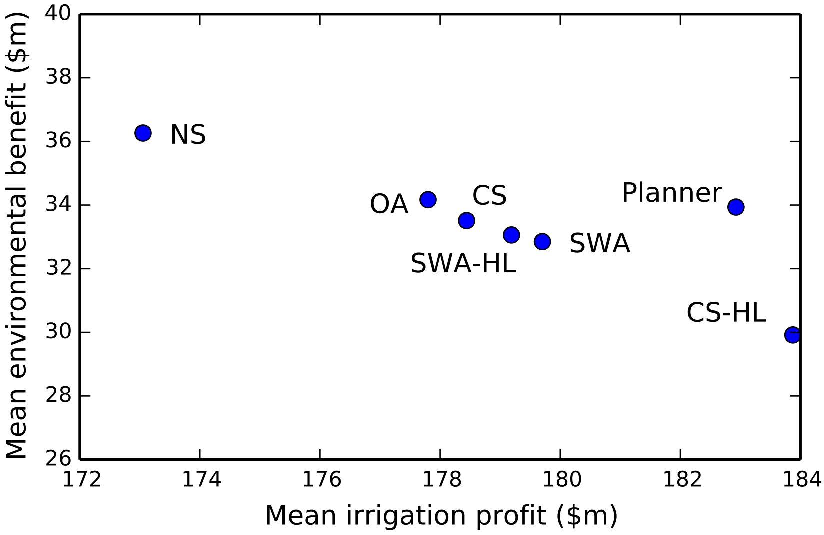

The mean profit and environmental benefit results are summarised in figure 7.10. Here we see a trade-off emerging between profits and the environment. The NS scenario leads to better environmental outcomes at the expense of profits, while CS-HL favours profits over environmental outcomes (for equal environmental shares).

A clearer picture of this 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 $[0.1, 0.5]$ (see appendix E). 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 7.11. The mean social welfare results are summarised in table 7.24.

Figure 7.11: Profit / environmental benefit trade-off curves, for $\lambda_{0t} = 0.1$ to 0.5

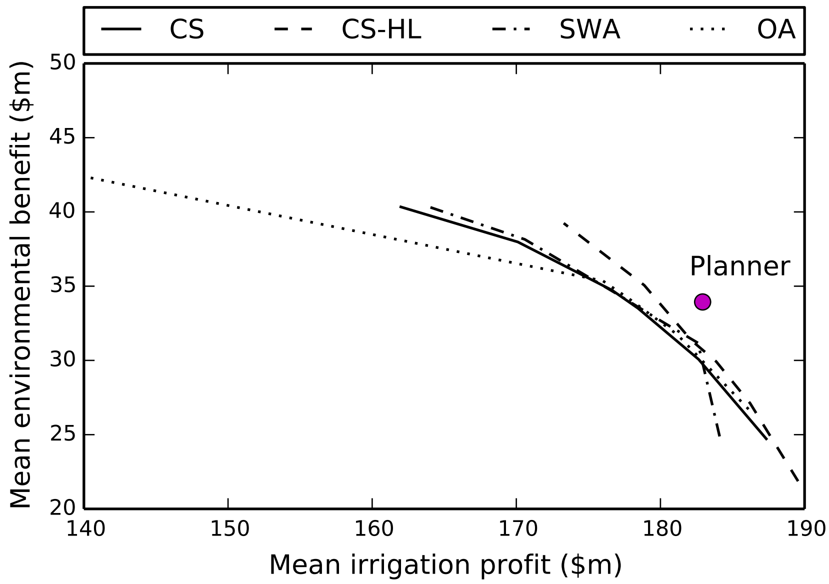

Table 7.24: Mean social welfare ($m) for $\lambda_{0t} \in [0.1, 0.2, 0.263, 0.3, 0.4, 0.5]$

	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 (see appendix E). 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 social welfare (table 7.24). 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. Once the reservoir is full all new inflows spill 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 7.8), it casts doubts over the suitability of OA for rivers with large EWHs.

Environmental water demand

Again environmental water demand q_{0t} is more variable than consumptive demand (tables 7.25, E.16 and E.17). Environmental demand is highest under the NS scenario. Under OA environmental demand is relatively low, however good river flows (and environmental benefits) are still achieved due to higher spills (table 7.23).

Table 7.25: 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

Table 7.26: Environmental allocation, a_{0t} (GL)

	Mean	SD	2.5th	25th	75th	97.5th
Planner	54.34	67.24	0.00	0.00	82.76	235.85
CS	126.30	103.14	10.79	50.55	175.69	381.84
SWA	101.37	70.04	6.18	43.00	157.16	243.95
OA	96.68	64.35	7.83	42.66	142.10	234.52
NS	153.33	103.06	0.00	69.37	225.69	358.58
CS-HL	102.57	92.41	0.00	19.57	179.86	278.03
SWA-HL	123.66	100.12	11.14	48.41	174.54	355.89

Figure 7.12 compares the river flow duration curves (at node 3) of the CS, CS-HL, OA and SWA scenarios. Here we see that the OA scenario — and to a lesser extent

Table 7.27: Environmental storage, s_{0t} (GL)

	Mean	SD	2.5th	25th	75th	97.5th
Planner	0.00	0.00	0.00	0.00	0.00	0.00
CS	114.19	61.71	22.37	67.87	152.67	247.03
SWA	142.41	65.82	42.88	89.57	191.13	283.52
OA	130.21	53.64	44.73	89.89	167.11	254.85
NS	110.26	65.01	15.56	57.61	154.67	247.03
CS-HL	102.89	76.47	0.00	40.52	158.96	247.03
SWA-HL	120.52	61.36	31.75	74.41	158.45	247.03

Figure 7.12: River flow duration curves, downstream F_{3t}

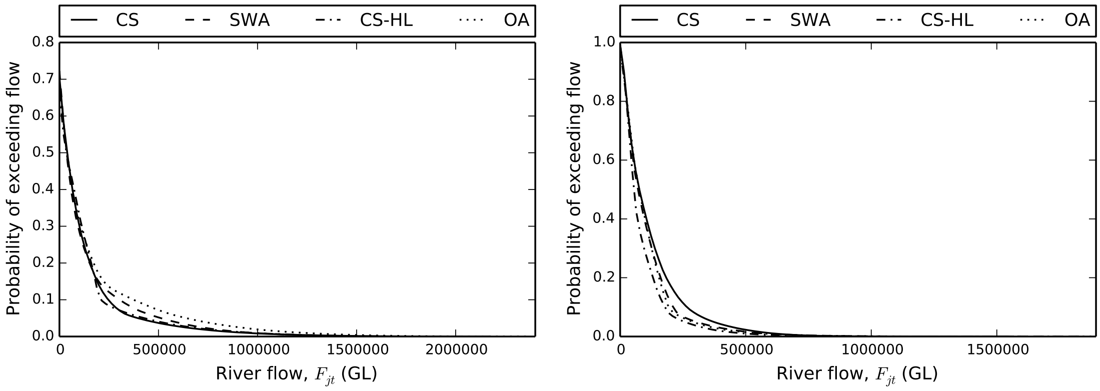

(a) Winter

(b) Summer

Figure 7.13: Mean environmental demand q_{0t} versus storage S_t and inflow I_t

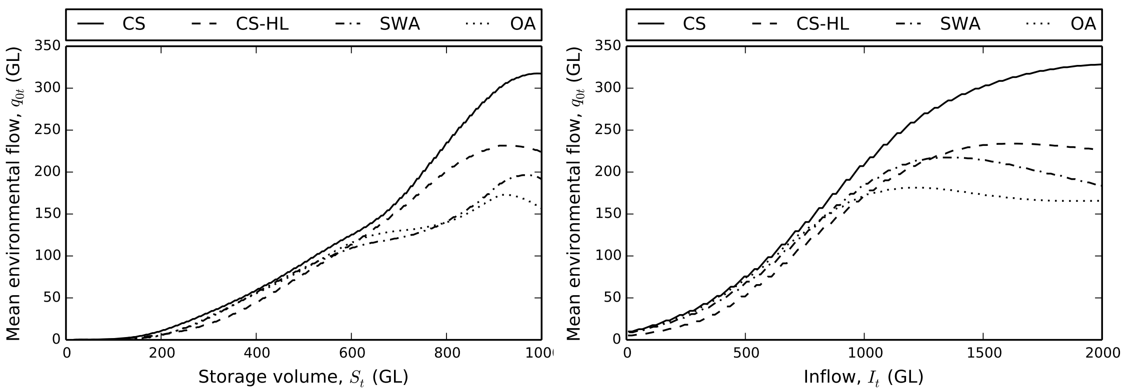

(a) versus storage

(b) versus inflow

the SWA scenario — result in more high flow events in winter (on account of spills), but lower river flows in summer.

Figure 7.14 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.

Environmental trade

Tables 7.28 to 7.30 show the EWH's 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 a net buyer in winter as would be expected. However, trading patterns vary considerably across years: in some years the EWH is a net buyer and in others a net seller.

Table 7.28: 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 7.29: 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 7.14 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

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

Table 7.30: 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

observed in previous studies, where the EWH is selling water during 'dry' periods and buying in 'wet'.

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.

Finally, CS-HL results in much less trade on average. The results suggest that low priority rights are a good match for the demands of the EWH, minimising their trade requirements.

Figure 7.14: Environmental trade $P_t(q_{0t} - a_{0t})$ versus storage S_t and inflow I_t

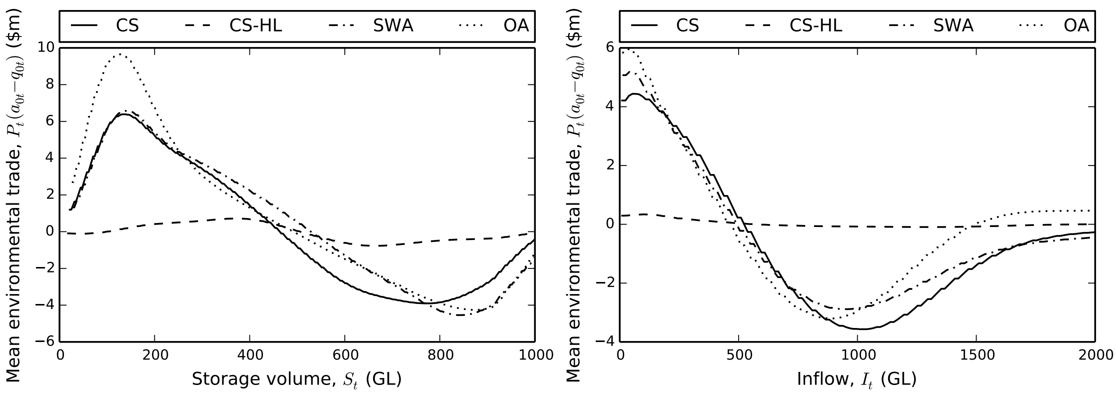

(a) versus storage

(b) versus inflow

In appendix E we present the results of a no-trade scenario and compute the gains from spot market trade (figure E.1). With environmental demands, the gains from trade are large (in the order of \$6m a year) even in scenarios with well defined storage rights and / or priority rights.

User group results

A selection of low and high reliability user group results are presented in tables E.14 to 7.26 in appendix E. The effect of the scenarios on the distribution of welfare (between low and high reliability user groups) is relatively trivial — in comparison with the trade-off between the environment and consumptive users in general.

7.7.2 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. As in chapters 5 and 6 we also present indexes relative to the CS scenario.

Social welfare

Mean social welfare results are summarised in tables 7.31-7.32 and figure 7.15. Here, CS-HL is the most frequently preferred scenario (198 of 547 complete runs), followed by OA (132), SWA (120) and CS (97).

Figure 7.15: Social welfare index, general case

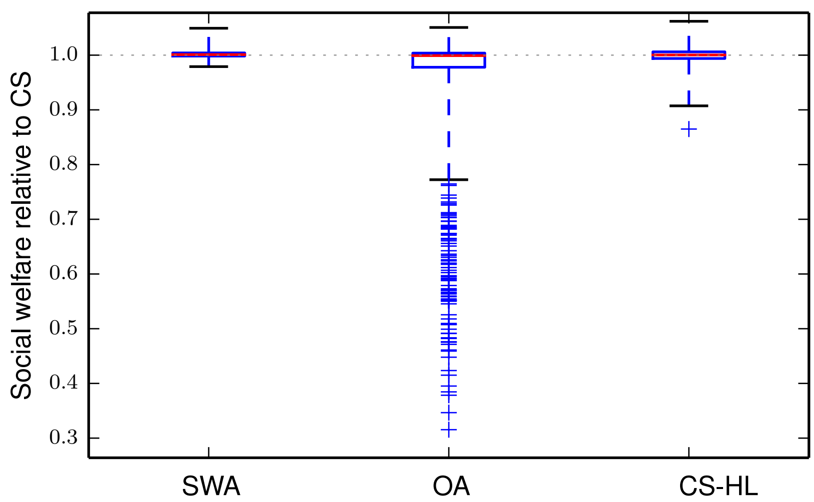

As in the central case, OA can lead to extremely low welfare under certain conditions. Figure 7.16 plots OA welfare relative to CS, against the environmental shares λ_0 and mean inflow relative to capacity $E[I_t]/K$. OA performs poorly where inflow is high relative to capacity (spills are frequent) and the environmental share is large.

Next we regress our mean welfare index against the model parameters. The most important parameters (table 7.33) are the environmental value ($b_\$/\bar{I}$), mean inflow

Table 7.31: 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 7.32: 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

Figure 7.16: OA welfare index, against $E[I_t]/K$ and λ_0

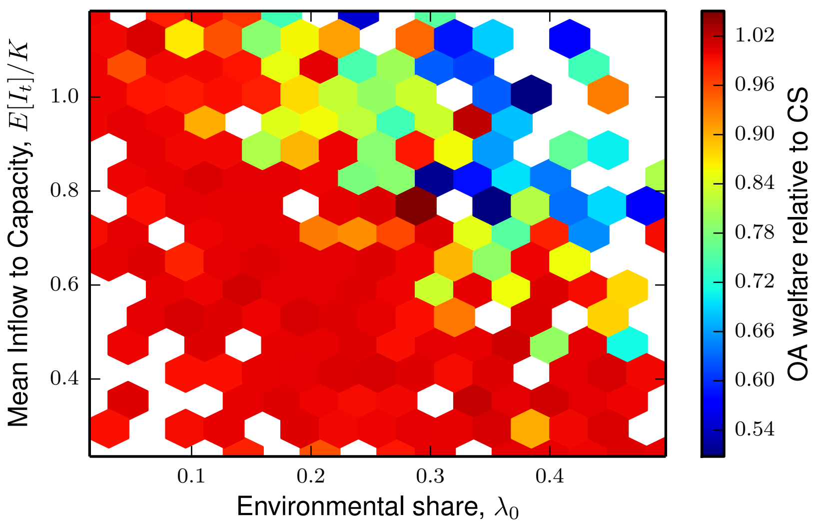

($E[I_t]/K$), the amount of high reliability demand (n_{high}) and the environmental inflow share 'shock' ($\lambda_0 - \hat{\lambda}_0$).

The effect of these parameters is summarised in figure 7.17. Together, $b_\$/\bar{I}$ and $\lambda_0 - \hat{\lambda}_0$ determine the environmental share λ_0. Figure 7.17 shows how the performance of OA quickly deteriorates for scenarios with high λ_0. OA also performs poorly for high values of $E[I_t]/K$ or n_{high}.

On average, CS-HL outperforms alternatives when $\lambda_0 > \hat{\lambda}_0$ (figure 7.17d). 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].

Figure 7.17: Social welfare index regression results

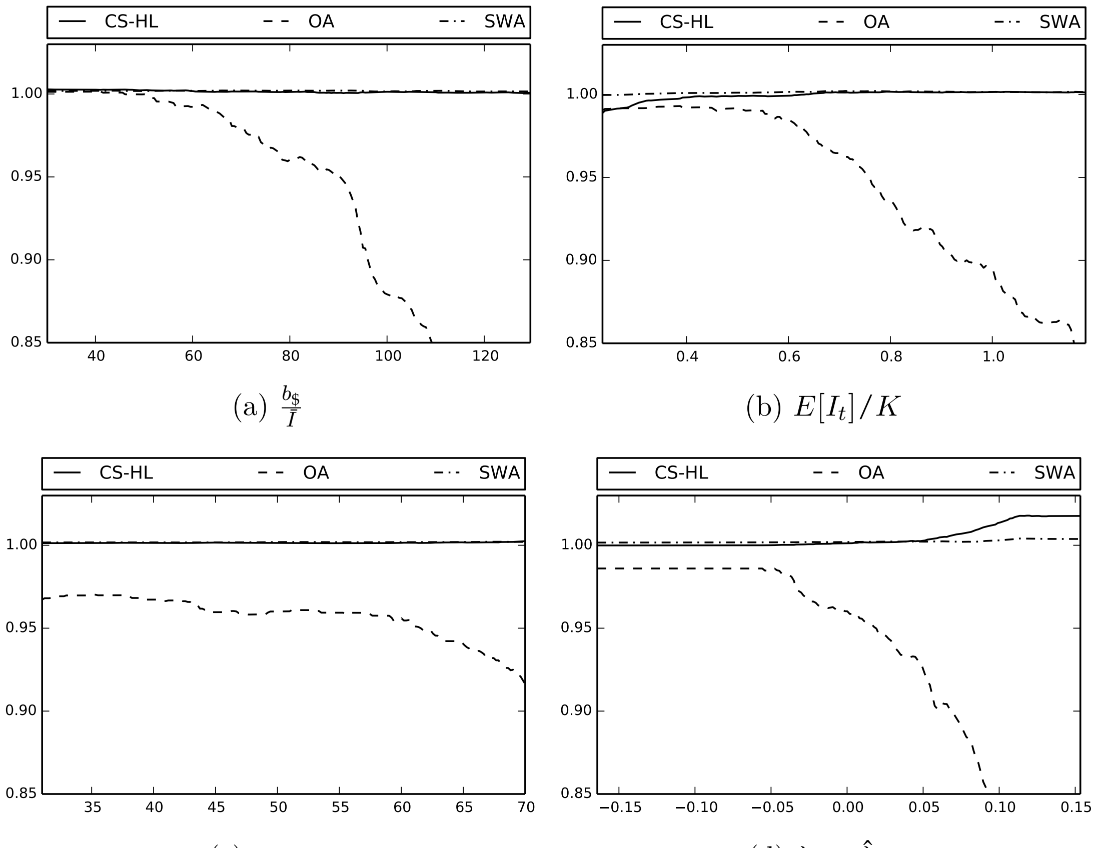

(a) $\frac{b_\$}{\bar{I}}$

(b) $E[I_t]/K$

(c) n_{high}

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

Next we regress the preferred scenario (as a qualitative dependent variable) against the parameters. Again the two most important parameters are $E[I_t]/K$ and $\lambda_0 - \hat{\lambda}_0$. Figure 7.18 plots the preferred scenario against $E[I_t]/K$ and $\lambda_0 - \hat{\lambda}_0$. Here the regression (classifier) model predicts CS-HL to be the preferred scenario, except for rivers with very low $E[I_t]/K$ and $\lambda_0 - \hat{\lambda}_0$ where OA is preferred.

[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 7.33: 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 7.18: Preferred scenario and classifier predictions by $E[I_t]/K$ and $\lambda_0 - \hat{\lambda}_0$

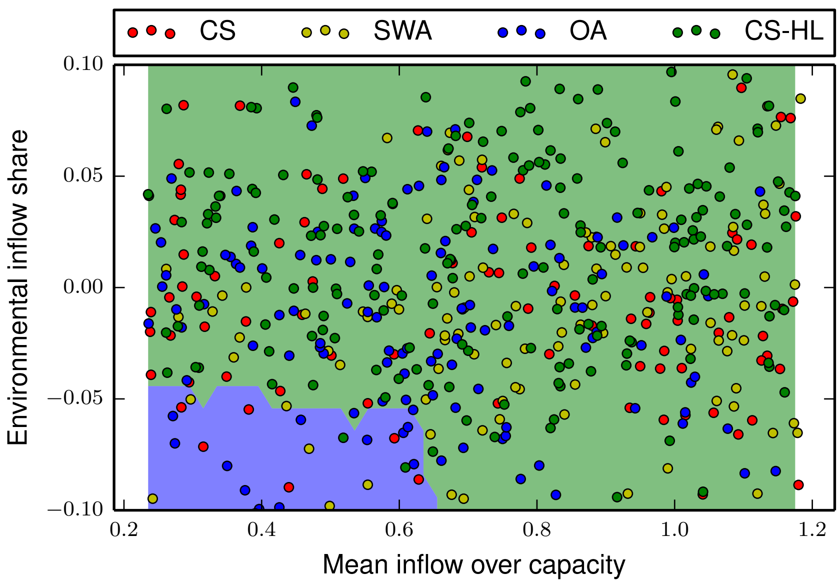

Table 7.34: Social welfare index classifier results

	Importance	CS	SWA	OA	CS-HL
$E[I]/K$	5.45	0.72	0.81	0.63	0.74
$\lambda_0 - \hat{\lambda}_0$	5.37	-0.01	-0.01	-0.01	0.02
δ_a	4.61	30,857.64	31,365.20	25,193.01	28,189.67
$\delta 0$	4.58	0.58	0.63	0.62	0.61
c_v	4.51	0.66	0.72	0.74	0.70
δ_R	4.47	0.12	0.10	0.10	0.11
$\frac{b_\$}{I}$	4.45	80.66	79.87	80.72	74.89
$\Lambda_{high}^{CS-HL} - \hat{\Lambda}_{high}^{CS-HL}$	4.45	0.00	0.00	0.00	0.01
σ_η	4.41	0.15	0.14	0.15	0.15
δ_{Ea}	4.34	34,609.91	41,473.51	33,818.33	35,471.47
α	4.30	9.28	8.87	8.70	9.36
b_1	4.23	0.31	0.31	0.32	0.32
ρ_e	4.20	0.41	0.41	0.40	0.40
σ_{e0}	4.15	0.51	0.47	0.48	0.52
μ_ω	4.13	0.64	0.65	0.65	0.64
ρ_I	4.12	0.25	0.25	0.25	0.25
τ	4.10	55.99	57.07	53.13	51.40
σ_ω	4.10	0.10	0.11	0.11	0.10
$\frac{\mathscr{A}_{low}}{E[I]/K}$	4.08	6,865.05	6,768.02	6,770.10	6,958.45
n_{high}	4.05	49.76	50.09	52.89	50.77
ω_δ	4.04	0.30	0.29	0.29	0.29
δ_{Eb}	3.96	0.21	0.20	0.20	0.20
$\Lambda_{high} - \hat{\Lambda}_{high}$	3.94	0.00	0.00	-0.00	0.00

Profit

Mean profits are summarised in tables 7.35-7.36 and figure 7.19. As in the central case, CS-HL tends to result in higher profits and OA in lower profits for a given environmental flow share (λ_0). Figure 7.19 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 7.20), storage levels (figure 7.21) and spills.

Figure 7.19: Mean profit index, general case

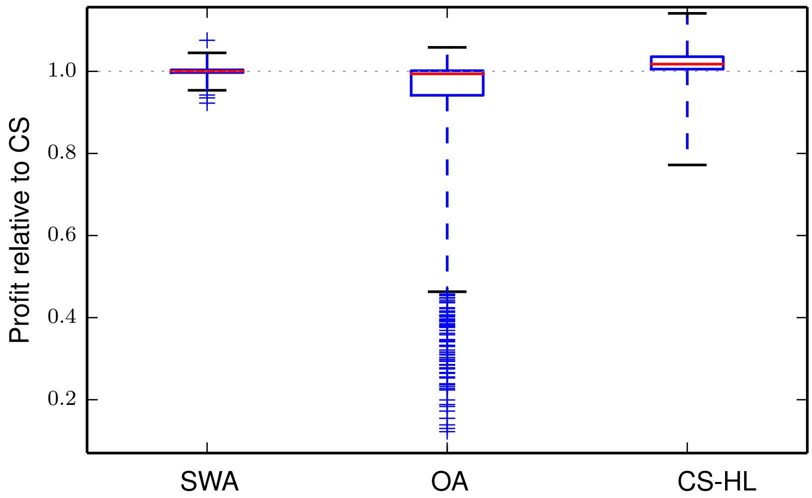

Table 7.35: Mean profit ($m), general case

	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 7.36: 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

Environmental benefits

Mean environmental benefits are summarised in tables 7.37-7.38 and figure 7.20. 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 7.20: Mean environmental benefit, general case

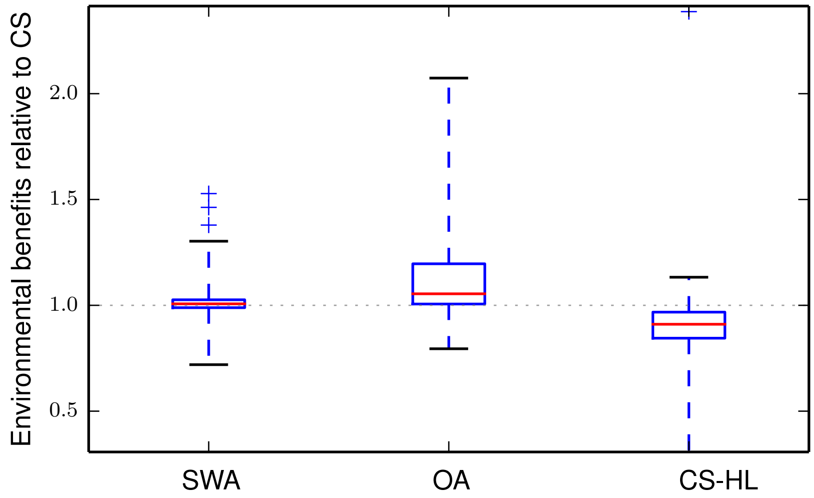

Table 7.37: Mean environmental benefits ($m), general case

	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 7.38: 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

Storage

Mean storage levels are summarised in tables 7.39 and 7.40 and figure 7.21. 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 7.21: Mean storage index, general case

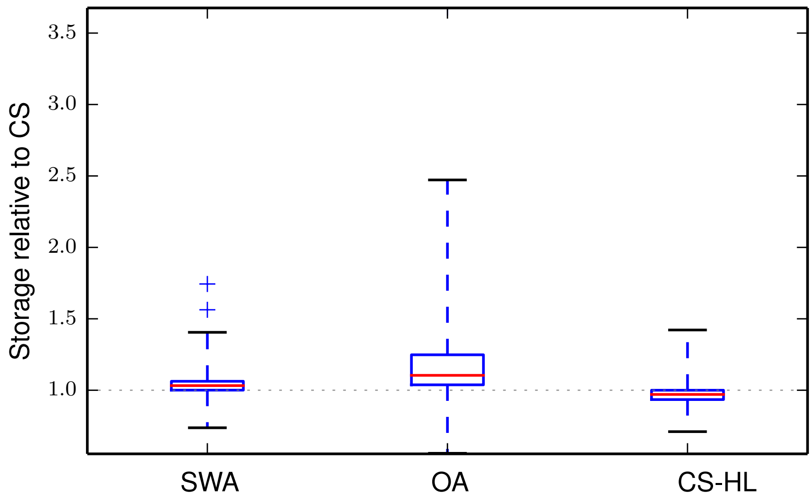

Table 7.39: Mean storage (GL), general case

	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 7.40: 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

7.8 Conclusions

In this chapter, we tested whether the introduction of a large EWH changed our conclusions regarding storage rights (chapter 5) and flow rights (chapter 6).

7.8.1 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. However, we observed similar relative effects to chapter 5: 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 be a disaster.

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.

7.8.2 Inflow rights

In chapter 6 we showed how storage rights 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 EWH 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.

7.8.3 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 present, most rivers in the MDB are converging on water rights systems which broadly resemble CS-HL (we discuss this further in chapter 9).

7.8.4 Future research

In this chapter, 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 (section 3.4.3) and a discretionary EWH. Another would be including flood mitigation objectives and rules.

While the focus of this chapter was on the design of water rights, the model could be used to consider EWH policy questions. For example, the optimal size and composition of the EWHs portfolio of rights and how this might vary with different types of environmental objectives or property rights.

Chapter 8

Solving large stochastic games with reinforcement learning

8.1 Introduction

Reinforcement learning (RL) provides a range of algorithms for solving Markov decision processes (MDPs), which do not require prior knowledge of the 'environment' (the payoff and transition functions). Rather agents 'learn' optimal polices by observing the outcomes of their actions (optimisation by simulation). Similar to dynamic programming, reinforcement learning works by exploiting the Bellman (1952) principle.

While common in artificial intelligence and operations research, reinforcement learning has received limited attention in economics[1]. One reason, is that for all their intuitive appeal, many practical challenges are faced in adapting these methods to economic problems. Developing solutions to these challenges has been an important part of this thesis.

We focus on the method of fitted Q iteration (Ernst et al. 2005); a batch version of standard Q-learning (Watkins and Dayan 1992). In fitted Q iteration, a large number of action, payoff and state transition samples are simulated, to which an action-value or Q function is then fit. Similar to fitted value iteration, the method is proven to converge, subject to assumptions on function approximation.

Ultimately, our goal is to develop methods to solve complex multi-agent problems, specifically stochastic games, where each agent faces a MDP with state transition and payoff functions dependent on the behaviour of other agents. Here we develop

[1]While historically related, the reinforcement learning methods sometimes applied in repeated games (see Erev and Roth 1998) are distinct from the methods we refer to here.

an approach in the spirit of 'multi-agent learning' (Fudenberg and Levine 2007), which combines reinforcement learning with game theory concepts.

Our approach provides a middle ground between the dynamic programming methods used in heterogeneous agent macro models and the simulation and search methods commonly used in agent-based computational economics. Both in terms of the size and complexity of the models it can be applied to, and the degree of rationality or 'intelligence' assumed for the agents.

We begin this chapter by defining our problem space: the single agent MDP and the stochastic game. We then provide an introduction to reinforcement learning for single agent problems, before detailing the function approximation techniques we employ, particularly tile coding. We then demonstrate the single agent method with an application to the planner's storage problem (chapter 3).

We then move on to multiple agent problems. Here we summarise the literature on solution concepts for stochastic games including Markov Perfect Equilibrium and Oblivious Equilibrium, before discussing the overlapping computer science and economic literature on multi-agent learning. We then introduce our multiple agent algorithm and detail its application to our water storage problems.

8.2 The problems

In this chapter we adopt some new notation. All notation is as defined in this section (except where we refer specifically to our water storage problems).

8.2.1 Markov decision process

A Markov decision process (MDP) represents the problem of an *agent* taking some *action* in an *environment* in order to maximise a *reward*. A MDP operates in discrete time: each period given current state s_t, the agent takes an action a_t, the environment then produces a reward r_t and a state transition s_{t+1}.

Formally, a Markov decision process is a tuple (S, A, T, R, β). $S \subset \mathbb{R}^{D_S}$ is the state space, where $D_S \in \{1, 2, ...\}$ is the dimensionality of the state space. $A \subset \mathbb{R}^{D_A}$ is the action space. $T : S \times A \times S \rightarrow [0, 1]$ is the transition function, a probability density function such that

$$\int_{S'} T(s, a, s')\ ds' = \text{Prob}(s_{t+1} \in S' | s_t = s, a_t = a)$$

$R : S \times A \to \mathbb{R}$ is the reward function. Finally $\beta \in (0, 1)$ is the discount rate. The agent's problem is to choose a_t to maximise the expected discounted reward

$$\max_{\{a_t\}_{t=0}^{\infty}} E\left\{\sum_{t=0}^{t=\infty} \beta^t R(a_t, s_t)\right\}$$

given T, R, s_0 and $a_t \in \Gamma(s_t) \subset A$, where Γ is the feasibility correspondence.

A (Markovian) *policy function* for the MDP is a mapping from states to actions $f : S \to A$. The discounted expected reward of following policy f is defined

$$V^f(s) = E\left\{\sum_{t=0}^{\infty} \beta^t R(f(s_t), s_t)|s_0 = s\right\}$$

where $s_{t+1} \sim T(s_t, f(s_t))$.

The *value function* associated with the optimal policy is defined as

$$V^*(s) = \sup_{f \in \Omega}\{V^f(s)\}$$

where Ω is the set of feasible policy functions.

Typically the value function also satisfies the Bellman principle

$$V^*(s) = \max_a \left\{R(s, a) + \beta \int_S T(s, a, s')V^*(s')\ ds'\right\}$$

8.2.2 Stochastic game

A stochastic game is essentially a multiple agent MDP.

There is a finite set of players $I = \{0, 1, ..., N\}$. The agents take actions $a_t^i \in A^i \subset \mathbb{R}^{D_A}$. We define the action profile as $a_t = (a_t^i)_{i \in I}$ and the action space $A = A^0 \times A^1 \times ... \times A^N$. The state of the game s_t, can include both agent specific states $s_t^i \in S^i \subset \mathbb{R}^{D_S}$ and a global state $s_t^g \in S^g \subset \mathbb{R}^{D_G}$, the state space is $S = S^0 \times S^1 \times ... \times S^N \times S^g$.

Each agent has reward function $R_i : S \times A \to \mathbb{R}$ and as before, we have a transition function $T : S \times A \times S \to [0, 1]$ and discount rate β. Then each agent i has to choose a_t^i to maximise their reward

$$\max_{\{a_t^i\}_{t=0}^{\infty}} E\left\{\sum_{t=0}^{t=\infty} \beta^t R_i(a_t, s_t)\right\}$$

given s_0, R_i, T, a_t^{-i} and $a_t^i \in \Gamma_i(s_t) \subset A^i$.

We define player policy functions $f_i : S \rightarrow A_i$ and the policy profile function $f : S \rightarrow A$. For any policy profile each player has a value function $V_i^f : S \rightarrow \mathbb{R}$ defined by

$$V_i^f(s) = E\left\{\sum_{t=0}^{\infty} \beta^t R((f(s_t), s_t)|s_0 = s\right\}$$

Stochastic games were first introduced by Shapley (1953), who showed that a two player zero-sum stochastic game could be solved by value iteration. Some of the main economic applications of stochastic games have been in industrial organisation, particularly models of oligopoly with investment and firm entry and exit (Ericson and Pakes 1995).

Stochastic games have also been applied to natural resource extraction problems (see Levhari and Mirman 1980). Recent applications have included fisheries (Kennedy 1987), groundwater (Negri 1989, Rubio and Casino 2001, Burt and Provencher 1993) and even surface water reservoirs (Ganji et al. 2007). Stochastic games have also been applied to commodity storage problems (Murphy et al. 1987, Rui and Miranda 1996).

8.3 Reinforcement learning

Reinforcement learning is a sub-field of machine learning concerned with solving MDPs. For a detailed introduction see Sutton and Barto (1998), Bertsekas and Tsitsiklis (1995) or Weiring and Otterlo (2012).

Figure 8.1: Reinforcement learning

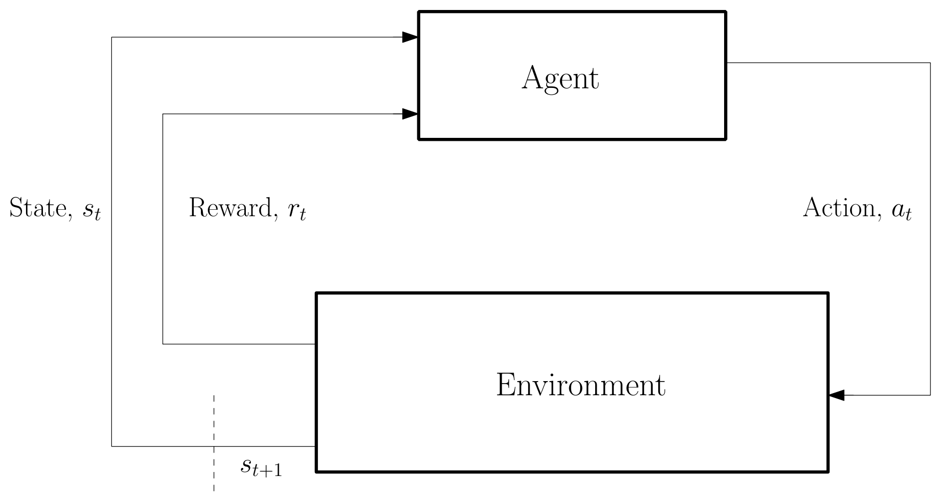

Central to reinforcement learning are so-called 'model free' approaches: where the transition and payoff functions are assumed unknown. Here the agent must learn only by observing the outcomes — the reward and state transition — of its interactions with the environment (figure 8.1). Learning a good policy then requires some 'exploration', that is testing a range of actions in each state.

Reinforcement learning can have some computational advantages over dynamic programming, particularly in larger problems. As a simulation method, attention is limited to realised state combinations (Judd et al. 2010), rather than a regular grid over the state space. Since, state variables are often correlated this can greatly reduce the complexity of the problem (Judd et al. 2010).

For example, figure 8.2 shows 10,000 simulated state points (I_t by S_t) for the planner's problem.

Figure 8.2: Planner's storage problem, sample state points

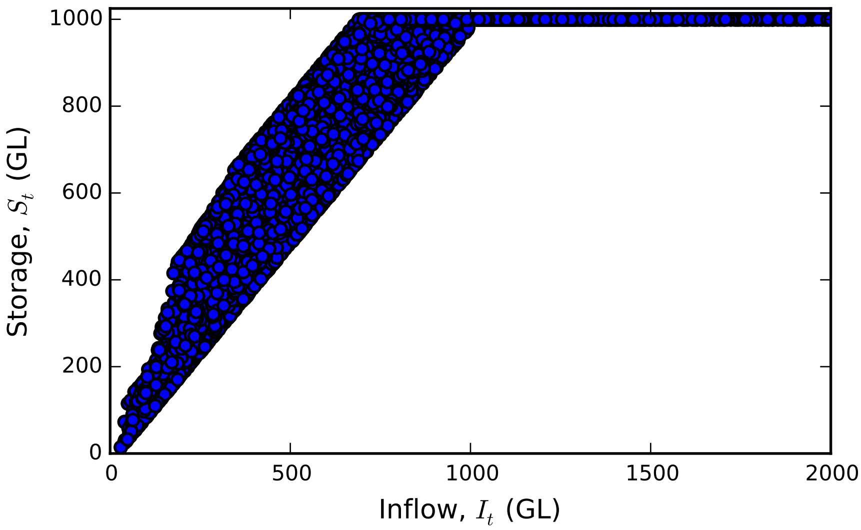

Many reinforcement learning methods are also easy to parallelize and generally provide greater flexibility to trade-off computation time and accuracy.

8.3.1 Q-learning

Q-learning (Watkins and Dayan 1992) is the canonical 'model free' reinforcement learning method. Q-learning works on the 'state-action' value function $Q : S \times A \rightarrow \mathbb{R}$, defined as the present value payoff from taking action a in state s and following an optimal policy thereafter

$$Q^*(a, s) = R(s, a) + \beta \int_S T(s, a, s') \max_a Q^*(a, s') \; ds'$$

Once in possession of Q^*, we can compute an optimal (aka greedy) policy without the payoff and transition functions

$$f^*(s) = \arg\max_a Q^*(a, s)$$

$$\max_a Q(a, s) = V^*(s)$$

In standard Q-learning we update the Q function after each state-action transition $\{s_t, a_t, r_t, s_{t+1}\}$. For a discrete state and action space, the algorithm operates as follows:

Algorithm 1: Q-learning with discrete state and actions

Initialise Q, s_0
for $t = 0$ *to* T **do**
  select $a_t \in A$; `// from some exploration policy`
  simulate (a_t, s_t) and observe r_t and s_{t+1}
  set $Q(a_t, s_t) = (1 - \alpha_t)Q(a_t, s_t) + \alpha_n\{r_t + \beta \max_a Q(a, s_{t+1})\}$
end

The actions a_t must be selected according to an 'exploration' (partially randomised) policy. The simplest option is an ϵ-greedy policy: a random policy with probability ϵ and an optimal policy otherwise. Here we encounter the standard exploration-exploitation trade-off: a highly random policy will provide good coverage of the state-action space, but risks spending too much time at irrelevant points.

$\alpha_t \in (0, 1)$ is known as the learning rate. α may be constant, but more commonly follows a decreasing schedule. Watkins and Dayan (1992) show (for discrete state and actions) that Q-learning converges as $t \to \infty$, subject to conditions over the exploration policy and learning rate α_t.

Q-learning can be extended to the continuous state and action case through function approximation. However, this typically voids convergence guarantees. More importantly, Q-learning is known to be unreliable (prone to spectacular divergence) in the continuous case (Weiring and Otterlo 2012).

8.3.2 Fitted Q iteration

Fitted Q iteration (Riedmiller 2005, Ernst et al. 2005) is a batch algorithm. First, a simulation is run for T periods (with an exploration policy). Then a series of Q function updates are applied to the batch of state-action samples (see algorithm 13).

The approach has two main advantages: *data efficiency* — since all samples are stored and reused — and *stability* — since Q functions can be fit to large samples. Further, it is well suited to multiple agent problems (Weiring and Otterlo 2012).

Algorithm 2: Fitted Q Iteration (continuous state and action)

initialise s_0
for $t = 0$ *to* T **do** `// Simulate the system for` T `periods`
    select $a_t \in A$; `// from some exploration policy`
    simulate (a_t, s_t)
    store the sample (s_t, a_t, r_t, s_{t+1})
end
initialise $Q(a_t, s_t)$
repeat `// Iterate until convergence`
    for $t = 0$ *to* T **do**
        set $\hat{Q}_t = r_t + \beta . \max_a .Q(a, s_{t+1})$
    end
    estimate Q by regressing $\hat{Q}_t$ against (a_t, s_t)
until *a stopping rule is satisfied*;

The separation of simulation and fitting stages permits much flexibility. Firstly, it allows any type of regression (supervised learning) model to be applied. Secondly, it facilitates parallel computing, since the simulation stage is so-called 'embarrassingly parallel'.

The success of fitted Q iteration depends crucially on function approximation. A variety of schemes have been proposed including random forests (Ernst et al. 2005), neural networks (Riedmiller 2005) and tile coding (Timmer and Riedmiller 2007). Similar to continuous dynamic programming, the algorithm is guaranteed to converge for 'non-expansive' type approximators (Ernst et al. 2005) (see section 8.4).

8.3.3 Fitted Q-V iteration

In noisy economic problems large samples T can be required. In this case, optimising Q for each state point s_t may be an unnecessary burden (especially in the multi-agent case). One option, is to optimise over a representative subset of state points, then estimate a continuous state-value function V (algorithm 3).

8.3.4 Sample grids

A natural choice for our subset of state points $\{s_k\}_{k=1}^{K}$ is a sample of approximately equidistant points (i.e., a sample grid). Our starting point here is a simple distance

Algorithm 3: Fitted Q-V iteration

initialise s_0
for $t = 0$ *to* T **do** `// Simulate the system for` T `periods`
  select $a_t \in A$; `// from some exploration policy`
  simulate (a_t, s_t)
  store the sample (s_t, a_t, r_t, s_{t+1})
end
initialise $Q(a_t, s_t)$
select a subset $\{s_k\}_{k=1}^{K} \subset \{s_t\}_{t=0}^{t=T}$
repeat `// Iterate until convergence`
  for $k = 0$ *to* K **do**
    set $\hat{V}_k = \max_{a_k} .Q(a_k, s_k)$
  end
  estimate $V(s_t)$ by regressing $\hat{V}_k$ on s_k
  for $t = 0$ *to* T **do**
    set $\hat{Q}_t = r_t + \beta.V(s_{t+1})$
  end
  estimate Q by regressing $\hat{Q}_t$ against (a_t, s_t)
until *a stopping rule is satisfied*;

based method (Algorithm 4), which provides a subset of points at least $\underline{r}$ distance apart[2]. This method is similar to the approach of Judd et al. (2010). However, our approach also counts the points within the radius $\underline{r}$ of each point: in order to identify outliers. Judd et al. (2010) remove outliers by separately estimating a density function.

Figure 8.3 provides a demonstration in two dimensions. This method is sufficient for moderate sample sizes but can become inefficient for very dense data sets. One option, is to add an early stopping condition, another is to employ some form of function approximation (i.e., tile coding) for more detail see appendix B.

[2]Note typically we scale all input data to the range [0, 1].

Algorithm 4: Selecting an approximately equidistant grid

set $J = 0$, $\mathbf{c}_0 = \mathbf{X}_0$, $n_0 = 0$
for $t = 0$ *to* T **do**
  set $r_{min} = \infty$
  for $j = 0$ *to* J **do** `// Find the nearest center` $\mathbf{c}_{j^*}$
    set $r = ||\mathbf{X}_t - \mathbf{c}_j||$
    if $r < r_{min}$ **then**
      set $r_{min} = r$
      set $j^* = j$
    end
  end
  if $r_{min} > \underline{r}$ **then** `// Add` $\mathbf{X}_t$ `as the next center` $\mathbf{c}_j$
    set $j = j + 1$
    set $\mathbf{c}_j = \mathbf{X}_t$
    set $n_j = 0$
  else `// Increment counter for center` $\mathbf{c}_{j^*}$ `by 1`
    set $n_{j^*} = n_{j*} + 1$
  end
end

Figure 8.3: An approximately equidistant grid in two dimensions

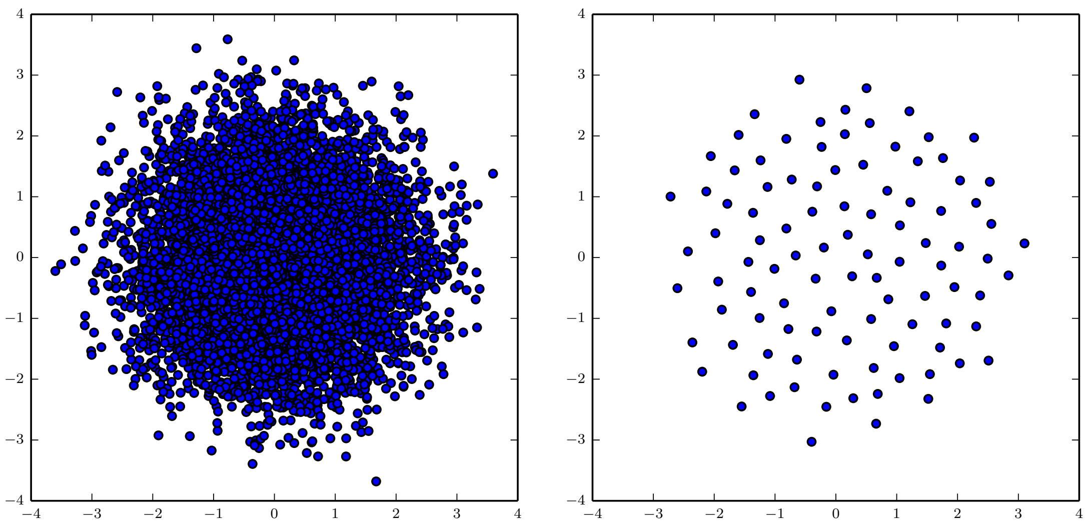

(a) 10,000 iid standard normal points

(b) 100 points at least 0.4 apart

8.4 Function approximation

The success of batch reinforcement learning depends crucially on function approximation. In machine learning this is known as 'supervised learning': to 'learn' (estimate) a model for a 'target' (dependent) variable Y, conditional a vector of 'input' (explanatory) variables $\mathbf{X}$, given only a set of 'training data' $\{Y_t, \mathbf{X}_t\}_{t=0}^{T}$. An example is provided in figure 8.4.

The goal with function approximation is prediction: we want a model that can accurately predict Y given $\mathbf{X}$ data outside our training sample. Minimising prediction error involves a bias-variance trade-off. A highly flexible model is at risk of 'overfitting' noisy data (figure 8.4a), while an inflexible model may lead to biased predictions (figures 8.4b, 8.4e).

For our purposes, computation time is also important: both fitting (estimation) time and prediction (function call) time. In practice, subtle trade-offs are faced between predictive power, fitting time and prediction time. Unfortunately, there is no general purpose method that achieves the optimal balance of all factors in all applications: there is no free lunch (Wolpert 1996).

Our fitted Q-V iteration approach, poses two distinct approximation problems: a *big problem* (i.e., the Q function) and a *small problem* (i.e., the policy and value functions f, V). These problems are summarised in table 8.1.

Table 8.1: Two approximation problems

	Big problem	Small problem
Sample size	Large (0.5-1 $\times 10^6$)	Small (500-2,000)
Input dimensions	Small (5)	Small (4)
Data structure	None	Gridded
Target noise	High	Low
Time constraint	Fitting	Prediction
Extrapolation important	No	Yes
Example	$Q(a_t, s_t)$	$f(s_t)$

For various reasons, standard methods employed in economics — such as orthogonal polynomials — are not ideal for these problems. Below, we introduce two methods common in reinforcement learning: radial basis function networks and tile coding. Ultimately, we used tile coding to approximate Q, V and f in all our water storage problems. Further detail on these methods is contained in appendix B.

Figure 8.4: A comparison of function approximation methods in one dimension

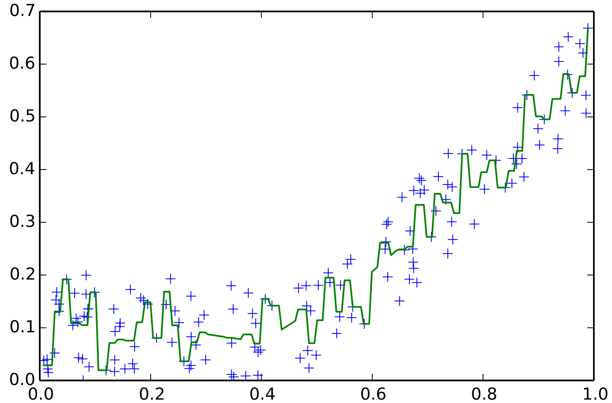

(a) Tile coding (TC-A) $N_L = 1, N_T = 60$

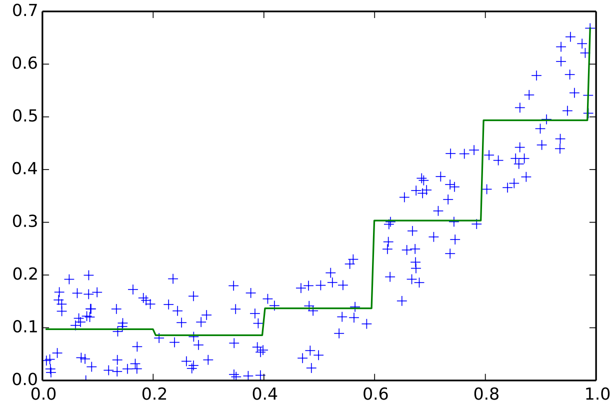

(b) Tile coding (TC-A) $N_L = 1, N_T = 6$

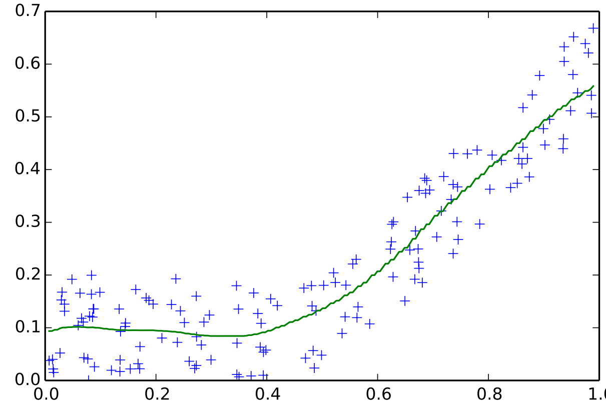

(c) Tile coding (TC-SGD) $N_L = 40, N_T = 4$

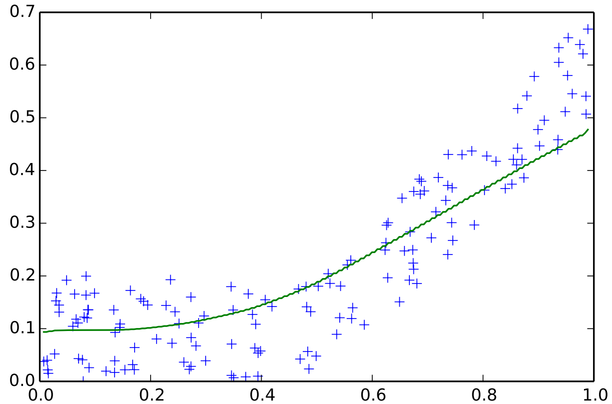

(d) Tile coding (TC-A) $N_L = 40, N_T = 3$

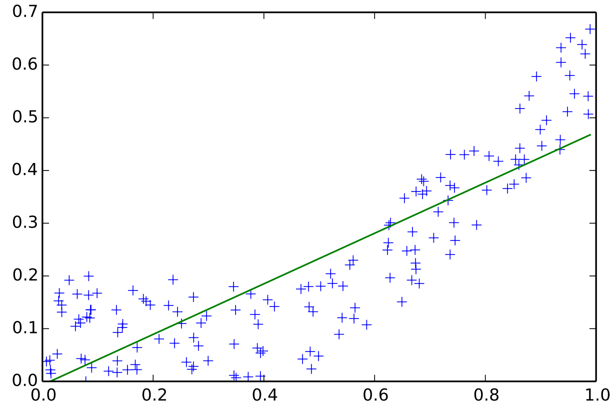

(e) Linear model

8.4.1 Radial basis function networks

A radial basis function (RBF) network is a linear combination of RBFs

$$\hat{Y}_t = \sum_{j=0}^{J} w_j \phi(||\mathbf{X}_t - \mathbf{c}_j||)$$

Here $\hat{Y}_t$ is our prediction of Y, w_j are parameters (weights) and ϕ is some RBF: a function of the euclidean distance r between an input vector and a fixed point (a center) c_j. Two common RBFs are the Gaussian and the thin-plate spline, respectively

$$\phi(r) = e^{-(\theta r)^2}$$
$$\phi(r) = r^2 \ln(r)$$

An RBF network is often combined with a low order polynomial function p. This allows the function to be flexible where we have data (and centers) and inflexible elsewhere: improving extrapolation.

$$\hat{Y}_t = \sum_{j=0}^{J} w_j \phi(||\mathbf{X}_t - \mathbf{c}_j||) + p(\mathbf{X}_t, \boldsymbol{\gamma})$$

With this type of RBF network we have up to four types of parameters to 'train': the centres $\mathbf{c}_j$, the RBF weights w_j, polynomial weights $\boldsymbol{\gamma}$ and in the gaussian case the 'bandwidth' θ^{-1}. A wide variety of estimation schemes exist, the preferred approach depends on the problem.

Small problems

For small problems, each sample point $\{\mathbf{X}_j\}_{t=0}^{t=T}$ can be a centre (an interpolation scheme). Alternatively, we can search for a good subset of points (i.e., a sparse model) via a model selection algorithm. A natural starting point for θ^{-1} is mean distance between input points. Given $\mathbf{c}_j$ and θ, w_j and $\boldsymbol{\gamma}$ can be fit by ordinary least squares (OLS).

Big problems

With large problems we need to select a small subset of points for the centres. A good option here is to use our sample grid method (algorithm 4) and then set $\theta^{-1} \approx \underline{r}$.

OLS is generally too slow in large data sets, so we turn to stochastic gradient descent (SGD) (Algorithm 5). Note that in all of these schemes $\{Y_t, \mathbf{X}_t\}_{t=0}^{t=T}$ are first scaled to $[0, 1]$.

Algorithm 5: Stochastic Gradient Descent: RBF network example

for $n = 0$ *to* N **do**
    for $t = 0$ *to* T **do**
        set $\delta_t = \sum_{j=1}^{j=J} w_j \phi(||\mathbf{X_t} - \mathbf{c}_j||) - Y_t$; `// prediction error`
        for $j = 0$ *to* J **do**
            set $w_j = w_j - \alpha_t \phi(||\mathbf{X_t} - \mathbf{c}_j||)\delta_t$; `// weight update`
        end
    end
end

Here the parameter $\alpha_t \in (0, 1)$ is known as the learning rate.

8.4.2 Tile coding

Tile coding (Albus 1975, Sutton and Barto 1998) is a function approximation scheme popular in reinforcement learning. With tile coding, the input space is partitioned into *tiles*. A whole partition is referred to as a *tiling* or a *layer*. A tile coding scheme then involves multiple overlapping layers, each offset from each other according to a *displacement vector*.

Figure 8.5: Tile coding

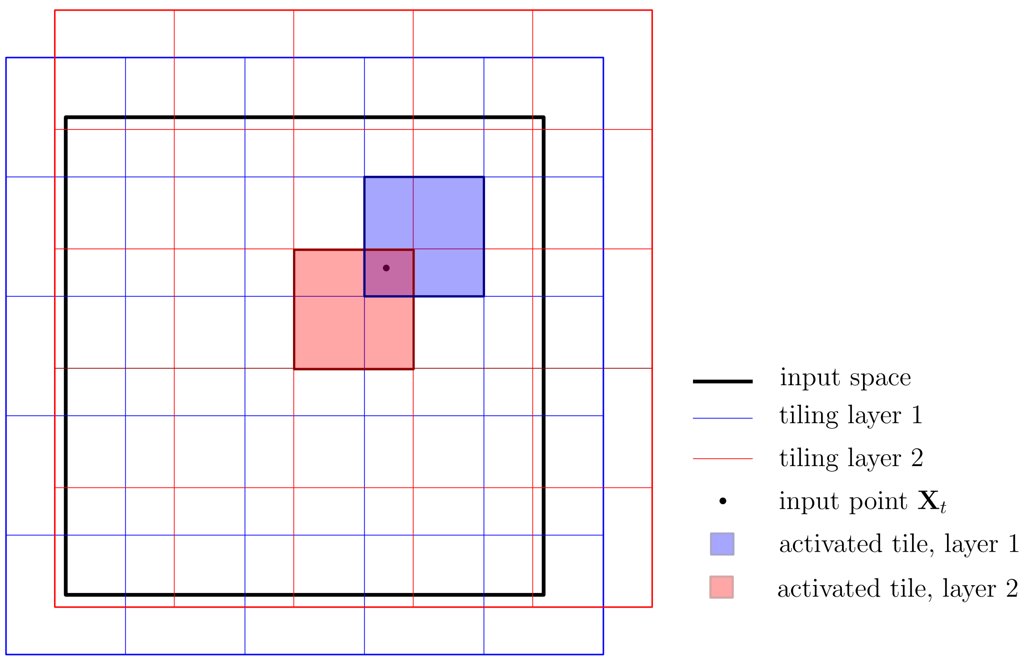

Tile coding is best understood visually. In the simplest approach the tilings are just regular grids (i.e. rectangular tiles) and each grid is offset uniformly (i.e. diagonally) as in figure 8.5. For any given point $\mathbf{X}_t$ one tile in each layer is *activated* and our predicted value is the mean of the weights attached to the active tiles.

More formally, a tile coding scheme involves $i = 1$ to N_L layers. Each layer contains $j = 1$ to N_T binary basis functions

$$\phi_{ij}(\mathbf{X}) = \begin{cases} 1 & \text{if } \mathbf{X} \in \mathscr{X}_{ij} \\ 0 & \text{if } \textit{otherwise} \end{cases}$$

where for each layer i the set $\{\mathscr{X}_{ij} : j = 1, ..., N_T\}$ is an exhaustive partition of the input space. The predicted value $\hat{Y}_t$ is then

$$\hat{Y}_t = \frac{1}{N_L} \sum_{i=1}^{N_L} \sum_{j=1}^{N_T} w_{ij} \phi_{ij}(\mathbf{X}_t)$$

where w_{ij} are the weights.

Tile coding is a constant piecewise approximation scheme: the 'resolution' of the approximation depends on the number of layers and the number of tiles per layer. For example, figure 8.4a shows a tile coding scheme where $N_L = 1$ and $N_T = 6$. Increasing the number of tiles gives a finer resolution but provides less *generalisation* leading to over fitting (figure 8.4b). Figure 8.4c, shows a model with $N_L = 40$ and $N_T = 4$ which provides both high resolution and good generalisation.

Tile coding has some computational advantages over RBF networks and other schemes with basis functions of global support. Tile weights are stored in arrays and accessed directly (computing array indexes simply involves integer conversion of $\mathbf{X}$, see appendix B). Each function call then involves only the N_L active weights. As such, prediction time is low and grows linearly in the number of layers rather than exponential in the number of dimensions.

Function predict($\mathbf{X}$)

set $\hat{Y} = 0$
for $i = 0$ *to* N_L **do**
    $j = index(\mathbf{X}, i)$;   `// returns index of active tile`
    $\hat{Y} = \hat{Y} + \frac{1}{N_L} w_{ij}$
end
return $\hat{Y}$

This speed gain comes at the cost of higher memory usage. However, since many reinforcement learning applications are CPU bound, this is an efficient use of resources. Memory limits are not an issue in any of our problems. In higher dimensions, weight arrays can be compressed through 'hashing' (see appendix B).

Fitting

The standard method of training the weights is SGD. An alternative, is to define each weight as a simple average, as in algorithm 6.

Algorithm 6: Fitting tile weights by averaging

for $t = 0$ *to* T **do**
    for $j = 0$ *to* N_L **do**
        $i = index(\mathbf{X}_t, j)$; `// returns index` i `of active tile`
        set $n_{ij} = n_{ij} + 1$; `// count tile sample size`
        set $w_{ij} = w_{ij} + Y_{ij}$; `// calculate sum`
    end
end
for $j = 0$ *to* N_l **do**
    for $i = 0$ *to* N_T **do**
        set $w_{ij} = w_{ij}/n_{ij}$; `// calculate mean`
    end
end

Timmer and Riedmiller (2007) considers tile coding in fitted-Q-iteration. When tile weights are simple averages, fitted-Q-iteration is guaranteed to converge on a unique fixed point (Timmer and Riedmiller 2007). A convergence result is possible, because this form of tile coding is a non-expansive approximator: that is a *smoother* or *averager* (Stachurski 2008, Gordon 1995)[3].

While fitting by averaging provides a convergence guarantee it is unlikely to provide ideal performance: it will suffer badly from bias if the tiles are too wide (see figure 8.4d) and variance if the tiles are too small. An alternative approach is averaged stochastic gradient descent (ASGD) (Bottou 2010), where the weights are defined as the average of a single SGD pass over the data (algorithm 7).

Bottou (2010) demonstrates the superiority of ASGD over SGD for problems with large samples. In section 8.5 we show that fitting the Q function by ASGD achieves a performance gain over averaging with no loss of stability.

[3] This form of tile coding is in fact closely related to other common averaging methods such as k-nearest neighbours and random forests.

Algorithm 7: Averaged Stochastic Gradient Descent (ASGD) — Tile coding

Initialise w_{ij} by averaging (algorithm 6)
set $\bar{w}_{ij} = 0$ for all i, j
for $t = 0$ *to* T **do** `// A single SGD pass`
  set $\delta_t =$ `predict` $(\mathbf{X})$ $-Y_t$
  for $j = 0$ *to* J **do**
    $i = index(\mathbf{X}_t, j)$
    set $w_{ij} = w_{ij} - \alpha_t \delta_t$
    set $\bar{w}_{ij} = \bar{w}_{ij} + w_{ij}$; `// sum the weight updates`
  end
end
for $j = 0$ *to* N_L **do**
  for $i = 0$ *to* N_L **do**
    set $w_{ij} = \bar{w}_{ij} / n_{ij}$; `// calculate mean`
  end
end

Big problems

For big problems (i.e., the Q function) we use tile coding with regularly spaced grids. We use the 'optimal' displacement vectors of (Brown and Harris 1994) (see appendix B). Tile weights are fit by ASGD.

Tile coding can suffer from noise in regions where training data are sparse, so for input variables with tails, we limit the tiling to a percentile range of the training data (e.g., the 1st to 99th). We then pass on the job of extrapolating into the unrepresented parts of the input space to our policy and value functions[4].

Small problems

For small problems we again use regular grids and optimal displacement vectors. The tile weights are fit either by averaging — for value functions — or standard SGD using averaging for starting values.

For extrapolation, we combine our tile coding scheme with a sparse linear spline model. The combined scheme replaces the tile code weight w_{ij} with the linear spline predicted value if $n_{ij} = 0$. For more detail see appendix B.

[4] There are a number of other more complex options here. One is the idea of 'adaptive tile coding' where the tile sizes are endogenous and may for example be larger in regions with less data (Whiteson et al. 2007). Another is to apply a non-linear scaling to the input data to make it more uniformly distributed.

8.5 The planner's problem

Here we apply fitted Q-V iteration to the planner's storage problem (section 3.5) and compare it to the benchmark of dynamic programming.

For fitted Q-V iteration we adopt a two stage 'growing batch' approach. We simulate an initial batch of samples assuming uniform (i.e., uninformed) exploration:

$$W_t = \epsilon_t . S_t$$

$$\epsilon_t \sim U[0, 1]$$

After computing an estimate $\hat{f}$ of the policy function from the first batch we add a second batch of samples, this time with Gaussian exploration:

$$W_t = \min\{\max\{\hat{f}(S_t, \tilde{I}_t) + \epsilon_t S_t, 0\}, S_t\}$$

$$\epsilon_t \sim N(0, \delta)$$

$$0 < \delta < 1$$

Algorithm 4 is used to build a grid of state points with $\underline{r} = 0.02$. Tile coding is used to approximate Q, V and f. We test fitting the Q function both by averaging (TC-A) with ASGD (TC-ASGD).

For dynamic programming we employ fitted policy iteration and use tile coding to approximate V over a 35×35 grid of the state space. We begin both methods with the initial guess $V(\mathbf{X}) = 0$.

Both methods are coded predominantly in Cython (see appendix B). Both make use of parallelization and run on a standard 4-core i7 desktop. Mean welfare, storage, and solution time are shown in tables 8.2-8.4 (each the average of 10 runs).

Fitted Q-V iteration obtains a policy comparable with SDP in less time. The fact that fitted Q-V iteration compares well for trivial single agent problems, suggests significant gains in computation time may be achievable in larger problems.

Reinforcement learning achieves welfare levels up to 99.8 per cent of the SDP solution and results in very similar mean storage levels.

For a given sample size ASGD outperforms averaging. Note that in this trivial example, TC-A still performs well on a computation time basis, because fitting time is longer with ASGD. However in more complex problems — where simulation is more time intensive — the gains from ASGD become more important.

Figure 8.6: Performance of fitted Q-V iteration, planner's storage problem

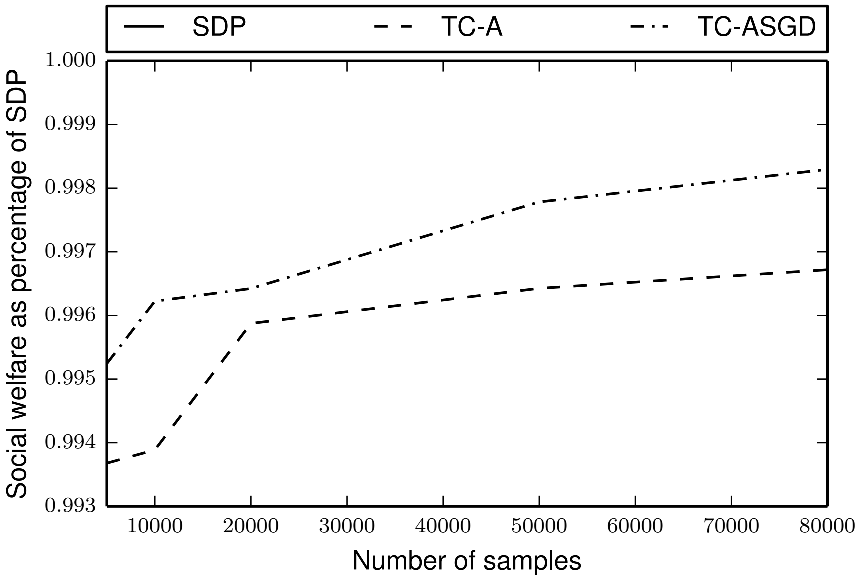

Table 8.2: Mean social welfare, by sample size

	5000	10000	20000	50000	80000
Myopic	181.4	181.4	181.4	181.4	181.4
SDP	186.6	186.6	186.6	186.6	186.6
TC-A	185.4	185.5	185.8	185.9	186.0
TC-ASGD	185.7	185.9	185.9	186.2	186.3

Table 8.3: Mean storage, by sample size

	5000	10000	20000	50000	80000
Myopic	577.4	577.4	577.4	577.4	577.4
SDP	697.8	697.8	697.8	697.8	697.8
TC-A	691.5	686.2	686.9	685.8	690.5
TC-ASGD	696.6	704.7	700.8	710.8	699.6

Table 8.4: Computation time, by sample size

	5000	10000	20000	50000	80000
SDP	6.6	7.2	7.5	7.4	7.4
TC-A	0.4	0.4	0.5	0.6	0.8
TC-ASGD	0.4	0.6	0.9	1.3	1.9

8.6 Multi-agent systems

8.6.1 Equilibrium concepts

Markov Perfect Equilibrium

A natural equilibrium concept for stochastic games is a Markov perfect equilibria (MPE) (Maskin and Tirole 1988). An MPE is defined by a set of markovian policies functions $\{f_0, f_1, ..., f_N\}$ which simultaneously solve each agent's problem, forming a sub-game perfect Nash equilibrium.

MPE existence results have been established for stochastic games as early as Shapley (1953). These early results typically assume finite state space or time horizons and mixed strategies. A general (pure strategy, infinite horizon and state space) existence result has remained elusive despite much recent attention (Duggan 2012, Escobar 2011, Balbus et al. 2011, Horst 2005, Amir 2005).

Broadly, MPE existence results involve two steps: one, show that for any feasible set of opponent policies f^{-i} the agent problems have unique solutions $V_i^*(s)$; two, show that the static 'stage game', with payoff functions $\pi_i(a, s, V_i(s))$

$$\pi_i(a, s, V_i(s)) = R(s, a) + \beta \int_S T(s, a, s')V_i(s')\ ds'$$

has a Nash equilibrium for any $s \in S$ and any feasible set of V_i.

Recent existence results all rely on particular regularizing assumptions. For example Escobar (2011) adopts an assumption of concave reduced payoffs: concavity of π_i with respect to a. Horst (2005) relies on a weak interaction condition: player's utility is affected more by their own actions than by others. Amir (2005) applies the lattice theory concepts of supermodularity and increasing differences. Recently Duggan (2012) proved the existence of MPE where the transition function is subject to a form of noise.

In general, uniqueness of equilibra in stochastic games is not guaranteed. As demonstrated by Doraszelski and Satterthwaite (2010) multiple equilibira are commonly observed in the Ericson and Pakes (1995) style models. The uniqueness of equilibria is usually considered numerically, by testing invariance to starting values. Although standard algorithms are not guaranteed to locate all possible equilibria (Borkovsky et al. 2008).

A general problem with stochastic games is that the state space ($S^0 \times S^1 \times ... \times S^N \times S^g$) can be too large, particularly for large N. Further, the assumption that the agents have information on all opponent state variables becomes unrealistic.

Oblivious equilibria (OE) is an alternative to MPE in the case where N is large (Weintraub et al. 2008). Here opponent state variables are replaced with relevant summary statistics.

Weintraub et al. (2008) show that — given a large number of similarly sized firms — OE approximates MPE for oligopoly type models. This result is generalised to a broader class of dynamic stochastic games by Abhishek et al. (2007). In the context of oligopoly models, opponent state variables are replaced with their long-run average means. Weintraub et al. (2010) extend OE to models with aggregate shocks, where opponent states are replaced with their mean conditional on the aggregate shock.

While uniqueness of OE is not guaranteed Weintraub et al. (2008) find no examples of multiple equilibria in applied problems, and argue that in general OE is likely to involve fewer equilibria than MPE.

8.6.2 Learning in games

The theory of learning in games describes how less than fully rational agents adapt in response to observed past play. There is much economic literature on learning in repeated games: testing how closely learning models reflect human behaviour in experiments (see Erev and Roth 1998) and establishing if and when they converge on equilibrium in models (see Fudenberg and Levine 1998).

Learning models are most relevant for games with large numbers of agents and 'aggregate statistics' where "players are only trying to learn their optimal strategy, and not to influence the future course of the overall system" (Fudenberg and Levine 2007; pp. 3). In the most general learning models, the population of agents can have 'heterogeneous beliefs', so that identical agents may play differing policies.

The oldest learning model is *fictitious play*: where each agent plays a best response to the empirical distribution of past play. A related model is the *partial best response* dynamic (Fudenberg and Levine 1998): where a sample of users play a best response to the previous periods play. Fudenberg and Levine (1998) show that for repeated games, these two models have identical asymptotic properties.

Another model popular in economics is 'reinforcement learning' — here we refer to the 'foresight-free' methods not the machine learning methods discussed previously. Here agents maintain a probability distribution over actions, with actions that result in higher payoffs, gradually receiving higher probabilities. Erev and Roth (1998) show that such simple rules closely match the behaviour of humans in experiments.

8.6.3 Multiple agent learning

Unfortunately, the economic literature on learning in stochastic games is surprisingly scarce (Fudenberg and Levine 1998). Here we turn to multi-agent learning: a relatively young but rapidly expanding field at the intersection of game theory and machine learning (Shoham et al. 2007, Fudenberg and Levine 2007, Busoniu et al. 2008). Here concepts of equilibrium and learning in repeated games meet reinforcement learning algorithms for single agent MDPs.

While reinforcement learning methods are designed for artificial 'software agents', they have a foundation in human and animal behavioural psychology and neuroscience (Weiring and Otterlo 2012). Putting aside this scientific 'inspiration', reinforcement learning provides a set of mature algorithms for representing agents who optimise subject to limited information and computational resources.

An obvious starting point for stochastic games, is to allow each agent to follow a single agent algorithm. In this case, the behaviour of the other agents becomes part of the environment that needs to be learned[5].

This type of multi-agent Q-learning has been applied widely in computer science domains, with some success (Busoniu et al. 2008). In the multi-agent context their are no convergence guarantees (as the environment is no longer stationary) and the convergence properties have been subject to limited study (Busoniu et al. 2008).

To date reinforcement learning has received little attention from economists:

> From the perspective of economists, Q-learning and other procedures that use generalizations of reinforcement learning to estimate value functions in environments with a state variable have not been well-studied [...] It may be that considering Q-learning in the multiple-agent case where players simultaneously try to calculate value function will lead to important new insights (Fudenberg and Levine 2007; pp. 6)

[5] In stochastic games incremental Q-learning with 'soft-max' exploration would be a natural analog to the Erev and Roth (1998) type methods used in repeated games.

> Surprisingly, the predictions of Q-learning models have yet to be compared with data from controlled laboratory experiments with human subjects — a good topic for future research (Tesfatsion and Judd 2006; pp. 979).

Economic applications of Q-learning are rare and the few examples (Tesauro and Kephart 2002, Kutschinski et al. 2003) apply Q-learning to repeated games (oligopoly price / quantity competition) rather than stochastic games.

In recent years, there has been a proliferation of more complex algorithms, that combine reinforcement learning methods with game theory concepts such as Nash equilibrium and fictitious play (see Busoniu et al. 2008). However, many of these methods have limited practical relevance as they are designed for narrow classes of games (e.g., zero-sum and two player games) and can be intractable in large problems (Busoniu et al. 2008, Shoham et al. 2007).

8.6.4 Multiple agent fitted Q-V iteration

In this thesis, we use a multi-agent version of fitted Q-V iteration (algorithm 8). In essence, the method combines our single agent algorithm with two smoothing dynamics (i.e., learning models). Similar to repeated games, a non-smoothed application of batch reinforcement learning will be unstable and prone to cycles.

One option is a *partial best response dynamic*. Within each simulation stage the environment is stationary (the users' policy functions are fixed) so we can compute optimal (best response) policies using fitted Q-V iteration, assign these to a random sample of the population, then generate a new sample and repeat.

The other obvious option is some form of fictitious play. Here all users would take the optimal polices at each stage, but new sample batches would be combined with existing samples (similar to the growing batch approach in section 8.5)

Our general multi-agent algorithm (algorithm 8) combines both types of smoothing:

This approach permits much flexibility. With high K and $\lambda = 1$ we have a partial best response dynamic, with low λ and K it approaches an 'on-line' reinforcement learning method. Our preferred approach (see section 8.7) is a comprise between these extremes.

This method is to be interpreted firstly as a learning algorithm. Within the computer science literature, this represents 'rational' agent learning (Bowling and Veloso 2001): learning that converges on best response policies given stationary

Algorithm 8: Multiple agent fitted *Q*-*V* iteration

```
simulate an initial batch of samples {s_t, a_t, r_t, s_{t+1}}_{t=0}^T
store the samples in arrays s, s_{+1}, a and r
initialise f_i, V_i, s_0, Q_i
for J iterations do
    for t = 0 to λT do                          // Simulate for λT periods
        select a_t using policies f_i with exploration
        simulate (a_t, s_t)
        store the new samples {s_t, a_t, r_t, s_{t+1}}
    end
    replace λT samples of s, s_{+1}, a and r with new samples
    select a subset of state points s̃ ⊂ s
    for K iterations do
        for i ∈ I do
            set Q̂^i = r^i + β.V^i(s_{+1})
            update Q_i, by regressing Q̂^i on (a^i, s)
            set V̂^i = max_a Q^i(a, s̃)
            update V_i by regressing V̂^i on s̃
        end
    end
    for i ∈ I^u do                              // update policy functions for I^u ⊂ I
        set â^i = arg max_a Q^i(a, s̃)
        update f_i by regressing â^i on s̃
    end
end
```

opponent policies. Within the economic literature, the approach might be described as an 'optimisation-based' learning method (Crawford 2013).

At the same time, the approach is only a small departure from algorithms used to compute rational expectations equilibria. For example, the approach is similar to those used to solve stochastic games for MPE, including the value iteration method of Shapley (1953) and the simulation based approach of Pakes and McGuire (2001). It is also related to the Krusell and Smith (1998) style algorithms used to solve macro heterogeneous agent models, where user problems are solved by dynamic programming and user interactions are estimated through simulation.

While our approach is clearly related to the field of Agent Based Computational Economics (ACE), there are some important differences. In practice, agent based models tend to rely on genetic algorithms / replicator dynamics and almost never make use of reinforcement learning methods (although they are considered in Tesfatsion and Judd 2006).

The differences between ACE and multi-agent learning reflect their respective origins. Agent based modelling comes from the physical sciences, with a focus on the aggregate dynamics resulting from very large numbers of simple agents. A typical example being animal herds. Multi agent learning comes from an engineering / computer science background. Here the focus is on how a moderate number of intelligent agents can develop good (optimal) policy rules. With typical examples being the interaction of robots / vehicles in logistics or defence.

8.7 The decentralised problem

Below we detail the application of fitted Q-V iteration to the decentralised storage problem of chapter 5.

We begin by solving the planner's problem by SDP. The planner's solution is then used to derive initial user policy functions. We then solve both the high and low reliability user problems (holding opponent policies fixed) by single agent fitted Q-V iteration. This yields an initial batch of T samples and estimates $\hat{f}_i, \hat{v}_i$ of the policy and value functions. We then proceed to the full multi-agent algorithm as outlined above.

We use $J = 25$ major iterations and $K = 1$ value iterations and $\lambda = 0.10$. After each major iteration we update policies for a 20 per cent random sample of agents. From this sample of agents we select a subsample to become 'explorers'. We adopt Gaussian exploration:

$$w_{it} = \min\{\max\{\hat{f}_i(s_{it}, S_t, e_{it}, \tilde{I}_t) + N(0, \delta_t.s_{it}), 0\}, s_{it}\}$$

$$0 < \delta_t < 1$$

The number of explorers and range of exploration declines over the 25 iterations: starting with 10 explorers (5 per user group) and $\delta = 0.25$ down to 4 explorers and $\delta = 0.085$.

We begin with T of 100,000, which gives us 500,000 samples for each of the user groups (given five explorers per group). We optimise the Q function for both groups over a sample grid of state points, with a radius of 0.045. Tile coding is used to approximate the policy and value functions and the Q function (fit by ASGD).

In chapter 5 we show that this method achieves a degree of convergence: while subject to some noise, changes in value and policy functions tend to diminish rather than cycle. Importantly, key model aggregate variables (prices, storage volumes and payoffs etc.) are stable. In testing, cycles tend to emerge in these problems if the degree of 'smoothing' is insufficient.

Further, we observe that scenarios involving very few externalities (e.g., CS), which are expected to achieve close to optimal outcomes, closely match the planner's SDP solution. Scenarios with large externalities (such as open access) result in expected changes in behaviour (over storage) and welfare losses.

Finally, if anything the multi-agent solution is observed to be more robust (i.e., achieves similar results when solved multiple times) than the single agent case. While individual policy functions tend to display some noise, this tends to average out across large numbers of agents.

8.8 Conclusions

Reinforcement learning provides a mature set of algorithms for solving MDPs. While some practical problems are faced in adapting them to economic problems, they are not insurmountable.

Economic problems tend to be noisy, relatively 'smooth' (the value and policy functions tend to be smooth after averaging out the noise) and involve relatively small state spaces. The batch method of fitted Q-V iteration is well suited to this context. While large samples may be required, the method can be faster than dynamic programming, when combined with appropriate function approximation.

Here, we find tile coding approximation to be ideal. In low dimensional problems, tile coding can process large data sets, much faster than alternatives relying on global basis functions. Tile coding may also be useful in other economic applications such as dynamic programming. In any more than 10 dimensions the memory requirements of tile coding will become restrictive[6].

For stochastic games, our multi-agent reinforcement learning methods provide a middle ground between Krusell and Smith (1998) style methods and agent based methods — both in terms of the size and complexity of the models it can be applied to and the degree of rationality or 'intelligence' assumed for the agents.

Reinforcement learning can handle problems too complex for dynamic programming based approaches. In particular, it allows us to consider decentralised economies with externalities. While, we may need to relax our notion of equilibrium in this case, we can hold tightly to the idea of individually maximising agents.

Clearly, the simulation and search methods of agent based economics provide maximum flexibility. However, agent based models tend to rely on simple behavioural rules (which may then replicate based on success). With reinforcement learning, we can have more 'intelligent' agent behaviour.

The field of multi-agent learning is still relatively young. There is much debate within and between disciplines on the best notion of equilibrium and how much emphasis to place on it (Fudenberg and Levine 2007). While, techniques are continuously evolving, the method of multi-agent fitted Q-V iteration, provides a practical starting point for economic problems.

[6]In large state spaces we can either turn to tree based methods such as 'random forests', or consider some form of state 'abstraction' / feature selection method to limit the number of state variables.

Chapter 9

Conclusion

9.1 Introduction

This thesis presented a series of computational experiments testing different water property rights systems. Our goal was to determine which system maximises allocative efficiency, given all are 'second best' policies subject to externalities and transaction costs.

Our focus was on regulated rivers: rivers controlled by large dams. Our experiments considered both property rights to river flows and to reservoir storage capacity.

To do this we developed a stylised model of a regulated river, in which a large number of water users make private trade and storage decisions. We then developed a novel 'artificial intelligence' algorithm, allowing us to populate the model with near optimal selfish agents.

We solved our model for a wide range of parameter values, with parameter ranges based on the Australian MDB. This allowed us to see if and how the preferred approach to water property rights varied with the nature of the river system.

Below we briefly summarise our findings. We then offer some remarks on: the methodology, the policy implications and directions for future research.

9.2 Findings

9.2.1 Storage rights

Chapter 5 considered storage rights in the absence of in-stream demands. The main alternatives were: the capacity sharing (CS) model of (Dudley and Musgrave

1988), spill forfeit rules (SWA) (Spillable Water Accounts), open access storage (OA) and no storage access (NS) (i.e., 'use it or lose it').

On average CS, was the highest welfare scenario, although the welfare differences between CS and SWA were mostly trivial. CS resulted in below optimal storage levels — on account of internal spills — and SWA above optimal. OA lead to substantial over-storage and NS substantial under-storage. As a result, both generally achieved lower welfare than CS and SWA. At a distributional level, over-storage (i.e., OA and SWA) favoured high reliability users.

Sensitivity analysis showed that the preferred approach to storage rights varied with the characteristics of the river system. The most important parameters were the mean and variance of inflow relative to storage capacity. OA performed relatively poorly in storage constrained rivers and relatively well (sometimes better than CS) when storage capacity was high relative to inflow.

We concluded that — in the absence of in-stream demands — either CS or SWA can achieve a good welfare outcome. Perhaps the most striking result was the magnitude of the storage effects. The results showed that small differences in water accounting rules can lead to large changes in user behaviour and in-turn aggregate storage volumes and river flows (via spills).

9.2.2 Flow rights

In chapter 6 we considered river flow rights, particularly priority versus proportional rights (again assuming no in-stream demand). Here, the main alternatives were capacity sharing with and without priority rights (CS, CS-HL) and release sharing — where storage decisions are made by a planner — with and without priority rights (RS, RS-HL).

With no trade, priority rights clearly outperformed proportional rights: CS-HL, RS-HL both outperformed CS, RS. However, the gains from priority rights decreased substantially (and were sometimes negative) in the presence of a spot market (even with a significant transaction cost). The introduction of risk aversion only weakened the performance of priority rights. The results also showed that storage rights have the ability to minimise trade requirements (e.g., CS outperformed RS in the absence of trade) further reducing the need for priority rights.

On average, the standard CS scenario outperformed the alternatives (even CS-U where inflow and capacity shares were 'unbundled'). The results supported the original CS model as specified in Dudley and Musgrave (1988). In the absence of in-stream demands, neither transaction costs nor risk aversion justifies the presence of priority rights. We concluded that priority rights are largely a hangover of the pre-trade era.

9.2.3 Environmental flows

Chapter 7 reconsidered storage and flow rights in the context of in-stream demands. Specifically, we introduced a large Environmental Water Holder (EWH) with payoffs defined over river flows. Here, the scenarios included CS, CS-HL, SWA, SWA-HL, NS and OA.

With in-stream values, storage rights generally led to under-storage relative to a planner's outcome: because most users ignore the environmental benefits of storage spills. Similar to chapter 5 we find that OA results in the highest storage levels, NS the lowest and SWA higher than CS.

The most striking result was the performance of the OA scenario. In some cases, OA performs well — achieving higher environmental benefits on account of storage spills — in others it is a disaster. Under certain conditions, OA leads the EWH to adopt a 'fill and spill' strategy: deliberately creating storage spills which benefit the environment, but block consumptive users from accessing water.

Overall, we found that CS-HL was the highest welfare scenario. This largely confirmed our findings on storage rights from chapter 5, but reversed our findings on priority rights from chapter 6. In the presence of a large EWH holder, the gains from spot market trading (between farmers and the EWH) can be large. Holding a large share of low reliability rights can significantly reduce the trade requirements and transaction cost exposure of the EWH.

9.3 Method

9.3.1 Reductive microeconomics

The results demonstrate the second-best nature of water markets (in the sense of Lipsey and Lancaster 1956) . We found many counter intuitive results, where seemingly sensible reforms — like internalising storage spills and losses, introducing priority rights, or unbundling storage and capacity shares — sometimes reduced social welfare[1]. Further, the results often varied with the nature of the river. For

[1]These findings are reminiscent of Brennan (2008*a*) who found that water trade could make things worse if storage rights were poorly defined.

example, OA often varied between the best and worst scenario depending on key parameters.

If water is so complex, and the results are so sensitive, one might ask to what extent can reductive microeconomics guide water policy?[2] Certainly the advocates of an institutional approach to water (chapter 4) seem to have their doubts.

In practice, water rights reforms are driven as much by user consultation and experimentation as they are by economic research. However, this trial and error process clearly has its costs. There are endless examples of water policy failures (see chapter 3) — a prime example being the recent failure of storage rights in northern Victoria (see chapter 5).

This thesis has shown that a microeconomic approach, when applied with discipline, can still have much predictive power in the area of water. Arguably, many of the results from our model correspond well with observation:

- OA and SWA type storage rights are only found in the southern MDB (where spills are infrequent) while capacity sharing type rights are found in the storage constrained northern regions (consistent with chapter 5).
- While we found no justification for priority rights — with trade and without an EWH — we know they evolved prior to trade (where the gains are significant).
- As inflow variation has increased (due to climate change) we have seen increased adoption of storage rights in the MDB (Hughes et al. 2013).
- Finally, the transition from a NS to a SWA / OA type system in northern VIC has seen significant increases in storage reserves (Hughes et al. 2013).

A related objection is that water institutions are too complex to be modelled. However, our model was easily able to accommodate institutions with complex water accounting rules, multiple private agents (i.e., irrigators) and government agencies (i.e., environmental water holders and reservoir managers). The approach could be extended to consider more complex systems, where decision makers are arranged in larger hierarchies.

While our policy scenarios remain somewhat abstract they still offer useful insights. Knowledge of the 'ideal' approach to water rights — even if it remains unattainable — can help guide reform in the right direction. Just as stylised models of water trade were useful in promoting the adoption of markets.

[2]A similar question was posed by Beare and Newby (2005; pp. 2): "The extent to which public resource policy can be guided by first principles remains an open question."

9.3.2 Computational experiments and reinforcement learning

Given the complexity of water, detailed models are required to achieve realistic results. Due to the second-best environment, analytical techniques are of little value. Proving that a policy is sub-optimal is not particularly useful, given the outcome can be anything from trivial to disastrous. Here the method of computational experiment (Kydland and Prescott 1996) has many obvious advantages.

Of course, modelling such complex multi-agent systems presents many challenges. In particular, the standard assumption of rational expectations can quickly become both intractable and unrealistic. If we limit ourselves to cases where a unique equilibrium can be confirmed we are potentially ruling out a large set of policy relevant models.

Agent-based modelling is a natural alternative, allowing researchers to develop more complex and realistic models. However, regardless of the fields greater intentions, most studies assign the agents simple behavioural rules and appeal to behavioural theories or results from human experiments.

Under a reinforcement learning approach, the agents gradually 'learn' near optimal policies. The method strikes a middle ground between rational expectations (i.e. dynamic programming) and agent-based methods. While it is based on simulation, computationally it is only a small departure from dynamic programming algorithms.

The relevance of equilibrium concepts in complex multi-agent problems remains an open question. The approach of this thesis — which follows the conventions of learning in games (Fudenberg and Levine 1998) — provides a pragmatic starting point. The design of solution methods and concepts for large stochastic games remains an important area of research in both computer science and economics.

9.3.3 Machine learning

This thesis introduced a number of machine learning techniques, both reinforcement learning and supervised learning. There is scope for much greater use of machine learning in economics and econometrics. Hal Varian recently made this point:

> I believe that these methods have a lot to offer and should be more widely known and used by economists. In fact, my standard advice to graduate students these days is "go to the computer science department and take a class in machine learning". (pp. 2 Varian 2014)

One of the strengths of machine learning, is that it frees the researcher to focus on the fundamentals of the problem. For example, it allows the researcher to focus on defining the constraints and payoffs (i.e., the simulation model) before selecting an appropriate learning algorithm to apply. This facilitates a more problem centred approach as opposed to a technique centred approach.

The machine learning methods introduced in this thesis are just the tip of the ice berg. Reinforcement learning, in particular, is a field of much diversity and progress. The next generation of algorithms will be more intelligent: requiring less human effort and 'tuning'. With improvements in computers and algorithms these methods will only become more powerful.

9.4 Policy

There has to be a layer of human judgment between model predictions and policy recommendations. Besides the obvious abstractions, there remain a number of factors that are inherently difficult to model including: ambiguity, history (i.e., path dependencies), politics, transaction / institution costs and departures from 'rational' behaviour. Taking these factors into account, our results largely support the policy recommendations for the MDB made previously by ABARES (Hughes 2014).

In terms of storage rights, the capacity sharing system of Dudley and Musgrave (1988) is a good template. At a high level, most regions in the MDB are converging on such an approach. Capacity sharing has already been adopted in some southern QLD rivers (Hughes and Goesch 2009*a*). While, the 'continuous accounting' systems of northern NSW represent an approximation of capacity sharing (Hughes 2014).

In the southern basin, storage rights loosely approximate capacity sharing in the sense that users have storage accounts and face account limits. However, for a number of reasons — particularly annual withdrawal accounting — these systems perform poorly in practice and can lead to open access type failures (see appendix C).

Hughes (2014) offers a number of policy recommendations, which would move MDB storage rights closer to the ideal of capacity sharing. For the most part, these are mundane adjustments to water accounting systems — the most notable being the introduction of 'continuous' (e.g., monthly or daily) withdrawal accounting in the southern basin. In all cases, the basic philosophy is that water accounting rules should reflect hydrological realities.

In chapter 6 we expressed concern with priority rights. We showed they offer little welfare gain in the presence of a spot market. Further, we argued they tend to exacerbate policy uncertainty and complicate share markets. This second problem is acute in the MDB given each region has its own priority rules resulting in a complex array of water products.

However, here the ideal of capacity sharing — homogeneous proportional shares in inflow — may be both unattainable (given path dependencies) and undesirable (given the results of chapter 7). Instead, there are some more minor reform options available (Hughes 2014) which could reduce policy uncertainty. In particular, the adoption of clear and consistent water allocation rules linking water supply (i.e., storage levels) and user allocations.

These property right reforms will become more important over time as climate change effects intensify and environmental water holders become more active.

9.5 Future research

Perhaps the most glaring omission from this study is flood mitigation. The inclusion of flood damage costs and a flood mitigation agency and / or rules is an obvious next step. The conflict between minimising exposure to drought on one hand and flood on the other is a growing issue in dam management, especially in the light of climate change.

This omission does not invalidate the results presented here though. Firstly, flood mitigation is typically a secondary concern in the MDB. Further, the storage rights systems considered here can easily operate in combination with flood mitigation rules: as they typically do in practice[3].

Other potential extensions include: allowing users to hold water right portfolios (see section 6.2.6), modelling different types of environmental objectives, modelling more complex river systems — with multiple storages, inflow sources and demand nodes[4] — and testing combinations of property rights and policy rules to see if and when mixed systems can outperform pure markets.

[3]A simple flood mitigation rule (see section 3.4.3) is equivalent to reducing the size of the storage from the perspective of water users. Here, the first problem is setting the right threshold level $\bar{S}$. The second issue is how different storage rights systems interact with flood mitigation rules, given they may alter the frequency with which the threshold binds.

[4]Here, an ultimate goal would be to combine our reinforcement learning type algorithm with existing hydrological models of actual river basins. This could become more feasible in time with improvements in algorithms and computing power. In the meantime, Hughes (2010) provides some thoughts on the implications of more complex river systems for capacity sharing type water rights.

Bibliography

Abhishek, V., Adlakha, S., Johari, R. and Weintraub, G. 2007, Oblivious equilibrium for general stochastic games with many players, *in* 'Allerton Conference on Communication, Control and Computing'.

Adamson, D., Mallawaarachchi, T. and Quiggin, J. 2006, State-contingent modelling of the Murray Darling Basin: implications for the design of property rights, Technical report, University of Queensland, School of Economics.

Adamson, D., Mallawaarachchi, T. and Quiggin, J. 2007, 'Water use and salinity in the Murray–Darling Basin: A state-contingent model', *Australian Journal of Agricultural and Resource Economics*, vol. 51, no. 3, pp. 263–281.

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.

Albus, J. S. 1975, 'A new approach to manipulator control: The cerebellar model articulation controller (CMAC)', *Journal of Dynamic Systems, Measurement, and Control*, vol. 97, no. 3, pp. 220–227.

Amir, R. 2005, Discounted Supermodular Stochastic Games: Theory and Applications.

Ancev, T. 2015, 'The role of the commonwealth environmental water holder in annual water allocation markets', *Australian Journal of Agricultural and Resource Economics*, vol. 59, no. 1, pp. 133–153.

Anderson, R. L. 1961, *The irrigation water rental market: a case study*, Economic Research Service, US Department of Agriculture.

Ansar, A., Flyvbjerg, B., Budzier, A. and Lunn, D. 2014, 'Should we build more large dams? The actual costs of hydropower megaproject development', *Energy Policy*.

Appels, D., Douglas, R., Dwyer, G. et al. 2005, 'Responsiveness of demand for irrigation water: A focus on the southern Murray-Darling Basin'.

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 Bureau of Agricultural and Resource Economics and Sciences 2012, *Drought in Australia: context, policy and management*, report to client edn.

Australian Bureau of Statisics 2012*a*, *Gross Value of Irrigated Agricultural Production, 2010-11*.

Australian Bureau of Statisics 2012*b*, 'Water account, Australia, 2010-11'.
<http://www.abs.gov.au/ausstats/abs@.nsf/mf/4610.0>

Australian Bureau of Statisics 2012*c*, 'Water use on Australian farms'.
<http://www.abs.gov.au/ausstats/abs@.nsf/mf/4618.0>

Australian Bureau of Statistics 2015, 'Australia's Environment: Issues and Trends'.
<http://www.abs.gov.au/ausstats/abs@.nsf/Lookup/4613.0main+features12007>

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

Baland, J.-M. and Platteau, J.-P. 1996, *Halting degradation of natural resources: is there a role for rural communities?*, Food & Agriculture Org.

Balbus, L., Reffet, K. and Wozny, L. 2011, A Constructive Study of Markov Equilibria in Stochastic Games with Strategic Complementarities, *in* 'Conference on Theoretical Economics, University of Kansas'.

Beare, S. 2010, 'Permanent Water Entitlements and Risk Management'.
<http://analytecon.com.au/permanent-water-entitlement.pdf>

Beare, S. C., Bell, R. and Fisher, B. S. 1998, 'Determining the Value of Water: The Role of Risk, Infrastructure Constraints, and Ownership', *American Journal of Agricultural Economics*, vol. 80, no. 5, pp. 916–940.

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'.

Beare, S. and Newby, J. 2005, Incomplete markets, excluded goods and natural resource management, *in* 'Australian Agricultural and Resource Economics Society Conference', ABARE.

Bell, R., Gali, J., Gretton, P. and Redmond, I. 2007, The responsiveness of Australian farm performance to changes in irrigation water use and trade, *in* 'Annual Conference of the Australian Agricultural and Resource Economics Society', Productivity Commission.

Bellman, R. 1952, 'On the theory of dynamic programming', *Proceedings of the National Academy of Sciences of the United States of America*, vol. 38, no. 8, pp. 716.

Bertsekas, D. and Tsitsiklis, J. 1995, Neuro-dynamic programming: an overview, *in* 'Decision and Control, 1995., Proceedings of the 34th IEEE Conference on', Vol. 1, IEEE, pp. 560–564.

Borkovsky, R., Doraszelski, U. and Kryukov, Y. 2008, 'A user's guide to solving dynamic stochastic games using the homotopy method'.

Bottou, L. 2010, Large-scale machine learning with stochastic gradient descent, *in* 'Proceedings of COMPSTAT'2010'.

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. 2006, 'Water policy reform in Australia: lessons from the Victorian seasonal water market', *Australian Journal of Agricultural and Resource Economics*, vol. 50, pp. 403–423.

Brennan, D. 2008*a*, '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. 2008*b*, 'Will we all be rooned without a desal plant? Hanrahan's lament and the problem of urban water planning under climate change', *Agenda*.

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.

Brown, M. and Harris, C. J. 1994, ‘Neurofuzzy adaptive modelling and control’.

Buchanan, J. M. 1962, ‘Externality’, *Economica*, vol. 29, no. 116, pp. 371–384.

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

Bureau of Rural Sciences 2007, *Science for decision makers: water banking.*

Burness, H. and Quirk, J. 1977, ‘Water rights and optimal reservoir management’.
<http://www.wordsmatter.caltech.edu/SSPapers/sswp165.pdf>

Burness, H. and Quirk, J. 1979, ‘Appropriative water rights and the efficient allocation of resources’, *The American Economic Review*, vol. pp. 25–37.

Burt, O. and Provencher, B. 1993, ‘The externalities associated with the common property exploitation of groundwater’, *Journal of Environmental Economics and Management*, vol. 24, no. 1, pp. 39–15K.

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.

Carey, J. M. and Sunding, D. L. 2001, ‘Emerging markets in water: a comparative institutional analysis of the Central Valley and Colorado-Big Thompson Projects’, *Nat. Resources J.*, vol. 41, pp. 283.

Cataldo, J. C., Behr, C., Montalto, F. A. and Pierce, R. J. 2010, ‘Prediction of Transmission Losses in Ephemeral Streams, Western U. S. A’, *Open Hydrology Journal*, vol. 4, pp. 19–34.

Challen, R. 2000, *Institutions, transaction costs, and environmental policy: institutional reform for water resources*, Edward Elgar Pub.

Cheung, S. 1998, ‘The Transaction Costs Paradigm 1998 Presidential Address Western Economic Association’, *Economic inquiry*, vol. 36, no. 4, pp. 514–521.

Chong, H. and Sunding, D. 2006, ‘Water Markets and Trading’, *Annual Review of Environment and Resources*, vol. 31, pp. 239–264.

Chong, J. and Ladson, A. R. 2003, ‘Analysis and management of unseasonal flooding in the Barmah–Millewa Forest, Australia’, *River Research and Applications*, vol. 19, no. 2, pp. 161–180.

Christensen, N. S., Wood, A. W., Voisin, N., Lettenmaier, D. P. and Palmer, R. N. 2004, 'The effects of climate change on the hydrology and water resources of the Colorado River basin', *Climatic Change*, vol. 62, no. 1, pp. 337–363.

Ciriacy-Wantrup, S. 1967, 'Water policy and economic optimizing: Some conceptual problems in water research', *The American Economic Review*, vol. 57, no. 2, pp. 179–189.

Commonwealth Scientific and Industrial Research Organisation 2008*a*, *Surface–groundwater connectivity assessment report-groundwater connectivity assessment report*.

Commonwealth Scientific and Industrial Research Organisation 2008*b*, 'The Murray-Darling Basin Sustainable Yields Project', http://www.csiro.au/en/Organisation-Structure/Flagships/Water-for-a-Healthy-Country-Flagship/Sustainable-Yields-Projects/MDBSY.aspx.

Commonwealth Scientific and Industrial Research Organisation 2010, *Climate variability and change in south-eastern Australia: A synthesis of findings from Phase 1 of the South Eastern Australian Climate Initiative*.

Commonwealth Scientific and Industrial Research Organisation 2012, *Assessment of the ecological and economic benefits of environmental water in the Murray-Darling Basin*.

Coumou, D. and Rahmstorf, S. 2012, 'A decade of weather extremes', *Nature Climate Change*, vol. 2, no. 7, pp. 491–496.

Crase, L. 2008, *Water policy in Australia: the impact of change and uncertainty*, Earthscan.

Crase, L., O'Keefe, S. and Dollery, B. 2009, 'Water Buy-Back in Australia: Policitic, Technical and Allocative Challenges'. <http://www.pc.gov.au/__data/assets/pdf_file/0020/90812/sub002.pdf>

Crase, L., Pagan, P. and Dollery, B. 2004, 'Water markets as a vehicle for reforming water resource allocation in the Murray-Darling Basin of Australia', *Water Resources Research*.

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.

Dahlman, C. J. 1979, 'The problem of externality', *Journal of law and economics*, vol. 22, no. 1, pp. 141–162.

Davidson, B. R. 1969, 'Australia wet or dry? The physical and economic limits to the expansion of irrigation.', *Australia wet or dry? The physical and economic limits to the expansion of irrigation.*

Davie, T. 2008, *Fundamentals of hydrology*, Taylor & Francis.

Davis, P. N. 1967, 'Australian and American water allocation systems compared', *Boston College Industrial and Commerical Law Review*, vol. 9, pp. 647.

Demsetz, H. 1967, 'Toward a theory of property rights', *The American Economic Review*, vol. 57, no. 2, pp. 347–359.

Donohew, Z. 2009, 'Property rights and western United States water markets', *Australian Journal of Agricultural and Resource Economics*, vol. 53, pp. 85–103.

Doraszelski, U. and Satterthwaite, M. 2010, 'Computable Markov-perfect industry dynamics', *The RAND Journal of Economics*, vol. 41, no. 2, pp. 215–243.

Draper, A. J., Jenkins, M. W., Kirby, K. W., Lund, J. R. and Howitt, R. E. 2003, 'Economic-engineering optimization for California water management', *Journal of water resources planning and management*, vol. 129, no. 3, pp. 155–164.

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

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*, '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 Burt, O. R. 1973, 'Stochastic reservoir management and system design for irrigation', *Water Resources Research*, vol. 9, no. 3, pp. 507–522.

Dudley, N. J. and Hearn, A. B. 1993, ‘Systems modeling to integrate river valley water supply and irrigation decision making under uncertainty’, *Agricultural systems*, vol. 42, no. 1, pp. 3–23.

Dudley, N. J., Howell, D. T. and Musgrave, W. F. 1971, ‘Optimal intraseasonal irrigation water allocation’, *Water Resources Research*, vol. 7, no. 4, pp. 770–788.

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.

Eggertsson, T. 1990, *Economic behavior and institutions: Principles of neoinstitutional economics*, Cambridge University Press.

Erev, I. and Roth, A. E. 1998, ‘Predicting how people play games: Reinforcement learning in experimental games with unique, mixed strategy equilibria’, *American economic review*, vol. pp. 848–881.

Ericson, R. and Pakes, A. 1995, ‘Markov-perfect industry dynamics: A framework for empirical work’, *The Review of Economic Studies*, vol. 62, no. 1, pp. 53–82.

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

Escobar, J. 2011, ‘Equilibrium Analysis of Dynamic Models of Imperfect Competition’.

Flinn, J. C. and Guise, J. W. 1970, ‘An application of spatial equilibrium analysis to water resource allocation’, *Water Resources Research*, vol. 6, no. 2, pp. 398–409.

Freebairn, J. and Quiggin, J. 2006, ‘Water rights for variable supplies’, *Australian Journal of Agricultural and Resource Economics*, vol. 50, pp. 295–312.

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.

Gale, M., Edwards, M., Wilson, L. and Greig, A. 2014, ‘The Boomerang Effect: A Case Study of the Murray-Darling Basin Plan’, *Australian Journal of Public Administration*, vol. 73, no. 2, pp. 153–163.

Ganji, A., Khalili, D. and Karamouz, M. 2007, ‘Development of stochastic dynamic Nash game model for reservoir operation. I. The symmetric stochastic model with perfect information’, *Advances in water resources*, vol. 30, no. 3, pp. 528–542.

Getches, D. 1984, *Water Law: in a nutshell*, West Publishing Company.

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

Gippel, C. 2006, ‘Scoping study of operational changes to reduce river and storage losses along the Eiver Murray system’, *Fluvial Systems Pty Ltd., Stockton, NSW. The Living Murray, Murray-Darling Basin Commission, Canberra.*

Glenn, E. P., Lee, C., Felger, R. and Zengel, S. 1996, ‘Effects of water management on the wetlands of the Colorado River delta, Mexico’, *Conservation Biology*, vol. 10, no. 4, pp. 1175–1186.

Gordon, G. J. 1995, Stable function approximation in dynamic programming, Technical report, DTIC Document.

Grafton, R. 2000, ‘Governance of the Commons: A Role for the State?’, *Land Economics*, vol. pp. 504–517.

Grafton, R. 2010, ‘How to increase the cost-effectiveness of water reform and environmental flows in the Murray-Darling Basin’, *Agenda: A Journal of Policy Analysis and Reform*, vol. 17, no. 2, pp. 17.

Grafton, R. and Jiang, Q. 2011, ‘Economic effects of water recovery on irrigated agriculture in the Murray-Darling Basin’, *Australian Journal of Agricultural and Resource Economics*.

Grafton, R., Landry, C., Libecap, G. and O’Brien, R. 2010, *Water Markets: Australia’s Murray-Darling Basin and the US Southwest*, National Bureau of Economic Research.

Grafton, R., Pittock, J., Davis, R., Williams, J., Fu, G., Warburton, M., Udall, B., McKenzie, R., Yu, X., Che, N. et al. 2012, ‘Global insights into water resources, climate change and governance’, *Nature Climate Change*.

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

Grafton, R. Q., Libecap, G., McGlennon, S., Landry, C. and OÂŠBrien, B. 2011*b*, 'An integrated assessment of water markets: a cross-country comparison', *Review of Environmental Economics and Policy*, vol. 5, no. 2, pp. 219–239.

Grant, R. M. 1996, 'Toward a knowledge-based theory of the firm', *Strategic management journal*, vol. 17, pp. 109–122.

Griffin, R. C. 2006, 'Achieving Water Use Efficiency in Irrigation Districts', *Journal of Water Resources Planning and Management-asce*.

Griffin, R. C. 2012, 'The Origins and Ideals of Water Resource Economics in the United States', *Annu. Rev. Resour. Econ.*, vol. 4, no. 1, pp. 353–377.

Griffin, R. C. and Characklis, G. W. 2002, 'ISSUES AND TRENDS IN TEXAS WATER MARKETING', *Water Resources Update*, vol. 121, pp. 22–33.

Griffin, R. C. and Hsu, S.-H. 1993, 'The Potential for Water Market Efficiency When Instream Flows Have Value', *American Journal of Agricultural Economics*.

Hafi, A., Thorpe, S. and Foster, A. 2009, The impact of climate change on the irrigated agricultural industries in the Murray–Darling Basin, *in* 'ABARE conference paper', Vol. 9.

Hanak, E. and Stryjewski, E. 2012, *California's water market, by the numbers: update 2012*, Public Policy Institute of California.

Hanemann, M. 2014, 'Property rights and sustainable irrigation: A developed world perspective', *Agricultural Water Management*, vol. 145, pp. 5–22.

Hardin, G. 1968, 'The tragedy of the commons', *New York*.

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'.

Heaney, A., Dwyer, G., Beare, S., Peterson, D. and Pechey, L. 2006, 'Third-party effects of water trading and potential policy responses', *Australian Journal of Agricultural and Resource Economics*, vol. 50, pp. 277–293.

Hirshleifer, J., De Haven, J. and Milliman, J. 1969, 'Water supply: economics, technology, and policy'.

Horst, U. 2005, 'Stationary equilibria in discounted stochastic games with weakly interacting players', *Games and Economic Behavior*, vol. 51, no. 1, pp. 83–108.

Howe, C., Alexander, P. and Moses, R. 1982, 'Performance of Appropriative Water Rights Systems in the Western United States during Drought, The', *Nat. Resources J.*, vol. 22, pp. 379.

Howe, C. W., Schurmeier, D. R. and Shaw, W. D. 1986, 'Innovative Approaches to Water Allocation: The Potential for Water Markets', *Water Resources Research*.

Howitt, R., Reynaud, A., Msangi, S. and Knapp, K. 2002, Calibrated Stochastic Dynamic Models For Resource Management, *in* '2002 Annual meeting, July 28-31, Long Beach, CA', American Agricultural Economics Association.

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. 2011, Estimating irrigation farm production functions using ABARES irrigation survey data, *in* 'Australian Agricultural and Resource Economics Society Conference'.

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

Hughes, N. 2014, 'Water storage rights: decentralising reservoir operation', *International Atlantic Economic Society Conference*.

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.

Hume, P. 2008, *Suggested method to estimate proportion of license needed to compensate irrigators for water transferred off irrigation schemes, a public submission on the draft water market rules to the ACCC*, Sustainable Soils Management.

Hurst, D. 2011, 'Three-quarters of Queensland a disaster zone'. <http://www.brisbanetimes.com.au/environment/weather/threequarters-of-queensland-a-disaster-zone-20110111-19mf8.html>

Iglesias, E., Garrido, A. and Gómez-Ramos, A. 2007, 'Economic drought management index to evaluate water institutionsÂŠ performance under uncertainty*', *Australian Journal of Agricultural and Resource Economics*, vol. 51, no. 1, pp. 17–38.

International Committee on Large Dams 2007, *Dams and the World's water*.

Johnson, R., Gisser, M. and Werner, M. 1981, 'The definition of a surface water right and transferability', *Journal of Law and Economics*, vol. 24, no. 2, pp. 273–288.

Judd, K. L., Maliar, L. and Maliar, S. 2010, A cluster-grid projection method: solving problems with high dimensionality, Technical report, National Bureau of Economic Research.

Katz, R. W., Parlange, M. B. and Naveau, P. 2002, 'Statistics of extremes in hydrology', *Advances in water resources*, vol. 25, no. 8, pp. 1287–1304.

Kennedy, J. 1987, 'A computable game theoretic approach to modelling competitive fishing', *Marine Resource Economics*, vol. 4, no. 1, pp. 1–14.

Kennedy, J. O. 1981, 'Applications of dynamic programming to agriculture, forestry and fisheries: Review and prognosis', *Review of Marketing and Agricultural Economics*, vol. 49, no. 3, pp. 141–173.

Kingsford, R. T., Walker, K. F., Lester, R. E., Young, W. J., Fairweather, P. G., Sammut, J. and Geddes, M. C. 2011, 'A Ramsar wetland in crisis–the Coorong, Lower Lakes and Murray Mouth, Australia', *Marine and Freshwater Research*, vol. 62, no. 3, pp. 255–265.

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.

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.

Kundzewicz, Z. W., Mata, L. J., Arnell, N. W., Döll, P., Jimenez, B., Miller, K., Oki, T., Şen, Z. and Shiklomanov, I. 2008, 'The implications of projected climate change for freshwater resources and their management', *Hydrological sciences journal*.

Kutschinski, E., Uthmann, T. and Polani, D. 2003, 'Learning competitive pricing strategies by multi-agent reinforcement learning', *Journal of Economic Dynamics and Control*, vol. 27, no. 11, pp. 2207–2218.

Kydland, F. E. and Prescott, E. C. 1996, 'The computational experiment: an econometric tool', *The Journal of Economic Perspectives*, vol. pp. 69–85.

Lane 1983, Transmission Losses, *in* 'National Engineering Handbook', United States Department of Agriculture, chapter 19.

Lefebvre, M., Gangadharan, L. and Thoyer, S. 2012, 'Do Security-Differentiated Water Rights Improve the Performance of Water Markets?', *American Journal of Agricultural Economics*, vol. 94, no. 5, pp. 1113–1135.

Levhari, D. and Mirman, L. 1980, 'The great fish war: an example using a dynamic Cournot-Nash solution', *The Bell Journal of Economics*, vol. pp. 322–334.

Libecap, G. 1978, 'Economic variables and the development of the law: the case of Western mineral rights', *Journal of Economic History*, vol. pp. 338–362.

Libecap, G. 2011*a*, 'Institutional path dependence in climate adaptation: Coman's some unsettled problems of irrigation', *American Economic Review*, vol. 101, no. 1, pp. 64–80.

Libecap, G. 2011*b*, 'Water rights and markets in the U.S. semi-arid west: efficiency and equity issues'.
<http://www2.bren.ucsb.edu/ glibecap/Lincolndevelopmentwaterrights.pdf>

Libecap, G. D. 2009, 'The tragedy of the commons: property rights and markets as solutions to resource and environmental problems', *Australian Journal of Agricultural and Resource Economics*, vol. 53, no. 1, pp. 129–144.

Lipsey, R. G. and Lancaster, K. 1956, 'The general theory of second best', *The review of economic studies*, vol. pp. 11–32.

Lund, J. R. 1996, *Developing seasonal and long-term reservoir system operation plans using HEC-PRM.*

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.

McCully, P. et al. 1996, *Silenced rivers: the ecology and politics of large dams.*, Zed Books.

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.

Murphy, F., Toman, M. and Weiss, H. 1987, 'A stochastic dynamic Nash game analysis of policies for managing the strategic petroleum reserves of consuming nations', *Management science*, vol. pp. 484–499.

Murray-Darling Basin Authority 2010, *Guide to the proposed Basin Plan*.

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

Murray-Darling Basin Authority 2012*a*, *Barmah-Millewa Forest Environmental Water Management Plan*.

Murray-Darling Basin Authority 2012*b*, *Proposed Basin Plan annual time series data*.
<http://www.mdba.gov.au/bpkid/>

Murray-Darling Basin Authority 2013, *The System for maintaining the River Murray Water Accounts*.

Murray-Darling Basin Authority 2014*a*, 'Annual Reports'.
<http://www.mdba.gov.au/about-mdba/corporate-documents/annual-reports>

Murray-Darling Basin Authority 2014*b*, 'Basin Plan'.
<http://www.mdba.gov.au/what-we-do/basin-plan>

Murray-Darling Basin Authority 2014*c*, 'Flood management at Hume dam'.
<http://www.mdba.gov.au/what-we-do/managing-rivers/river-murray-system/dams-weirs/hume-dam/flood-management-at-hume-dam>

Murray-Darling Basin Authority 2014*d*, 'History of the Basin Plan'.
<http://www.mdba.gov.au/what-we-do/basin-plan/development/history>

Murray-Darling Basin Authority 2014*e*, 'Sustainable Diversion Limits'.
<http://www.mdba.gov.au/what-we-do/water-planning/sdl>

Murray-Darling Basin Authority 2015, 'Live river data'.
<http://www.mdba.gov.au/river-data/live-river-data>

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

National Aeronautics and Space Administration 2013, 'NASA Study Projects Warming-Driven Changes in Global Rainfall'.
<http://www.nasa.gov/home/hqnews/2013/may/HQ_13-119_Rainfall_Response.html>

National Water Comission 2004, *Intergovernmental agreement on a national water initiative.*
<http://www.nwc.gov.au/reform/nwi>

National Water Comission 2010, *The impacs of water trading in the Southern Murray-Darling Basin.*

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

National Water Comission 2011*b*, *Australian water markets: trends and drivers 2007-08 to 2010-11.*

National Water Comission 2011*c*, *Water markets in Australia: A short history.*

Negri, D. 1989, 'The common property aquifer as a differential game', *Water Resources Research*, vol. 25, no. 1, pp. 9–15.

New South Wales Irrigation Council 2014, 'Entitlement conversion policy'.
<http://www.nswic.org.au/pdf/policy_documents/>

North, D. C. 1990, *Institutions, institutional change and economic performance*, Cambridge university press.

Oakland, W. H. 1974, 'Public goods, perfect competition, and underproduction', *The Journal of Political Economy*, vol. 82, no. 5, pp. 927–939.

Oates, W. E. 1999, 'An essay on fiscal federalism', *Journal of economic literature*, vol. 37, pp. 1120–1149.

Ostrom, E. 2010*a*, 'Beyond markets and states: polycentric governance of complex economic systems', *The American economic review*, vol. 100, no. 3, pp. 641–672.

Ostrom, E. 2010*b*, 'Polycentric systems for coping with collective action and global environmental change', *Global Environmental Change*, vol. 20, no. 4, pp. 550–557.

Ostrom, E., Gardner, R. and Walker, J. 1994, *Rules, games, and common-pool resources*, University of Michigan Press.

Ostrom, V. 1962, ‘The political economy of water development’, *The American Economic Review*, vol. 52, no. 2, pp. 450–458.

Ostrom, V. and Ostrom, E. 1972, ‘Legal and Political Conditions of Water Resource Development.’, *Land economics*.

Pakes, A. and McGuire, P. 2001, ‘Stochastic algorithms, symmetric Markov perfect equilibrium, and the ÂŚcurseÂŠof dimensionality’, *Econometrica*, vol. 69, no. 5, pp. 1261–1281.

Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V., Vanderplas, J., Passos, A., Cournapeau, D., Brucher, M., Perrot, M. and Duchesnay, E. 2011, ‘Scikit-learn: Machine Learning in Python’, *Journal of Machine Learning Research*, vol. 12, pp. 2825–2830.

Peel, M. C., McMahon, T. A. and Finlayson, B. L. 2010, ‘Vegetation impact on mean annual evapotranspiration at a global catchment scale’, *Water Resources Research*, vol. 46, no. 9, pp. W09508.

Perera, B. and Codner, G. 1988, ‘A combined stochastic dynamic programming-statistical disaggregation approach applied to multiple reservoir systems’, *Water Resources Management*, vol. 2, no. 3, pp. 153–171.

Perera, B. and Codner, G. P. 1996, ‘Reservoir targets for urban water supply systems’, *Journal of Water Resources Planning and Management*, vol. 122, no. 4, pp. 270–279.

Poff, N. and Zimmerman, J. 2010, ‘Ecological responses to altered flow regimes: a literature review to inform the science and management of environmental flows’, *Freshwater Biology*, vol. 55, no. 1, pp. 194–205.

Productivity Comission 2010, *Market mechanisims for recovering water in the Murray-Darling Basin*.

Quiggin, J. 1993, ‘Common property, equality, and development’, *World Development*, vol. 21, no. 7, pp. 1123–1138.

Quiggin, J. 2012, ‘Murray-Darling plan: a disastrous process, a sliver of hope’. <http://www.crikey.com.au/2012/05/31/murray-darling-plan-a-disastrous-process-a-sliver-of-hope/>

Radner, R. 1992, ‘Hierarchy: The economics of managing’, *Journal of economic literature*, vol. pp. 1382–1415.

Randall, A. 1974, ‘Coasian externality theory in a policy context’, *Nat. Resources J.*, vol. 14, pp. 35.

Randall, A. 1981, ‘Property entitlements and pricing policies for a maturing water economy’, *Australian Journal of Agricultural Economics*.

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

Rani, D. and Moreira, M. M. 2010, ‘Simulation–optimization modeling: a survey and potential application in reservoir systems operation’, *Water resources management*, vol. 24, no. 6, pp. 1107–1138.

Riedmiller, M. 2005, Neural fitted Q iteration–first experiences with a data efficient neural reinforcement learning method, *in* ‘Machine Learning: ECML 2005’, Springer, pp. 317–328.

Ross, A. 2012, Integrated surface water and groundwater management in the Murray-Darling Basin, Colorado and Idaho, PhD thesis, Australian National University.

Rubio, S. and Casino, B. 2001, ‘Competitive versus efficient extraction of a common property resource: The groundwater case’, *Journal of Economic Dynamics and Control*, vol. 25, no. 8, pp. 1117–1137.

Rui, X. and Miranda, M. 1996, ‘Solving nonlinear dynamic games via orthogonal collocation: An application to international commodity markets’, *Annals of Operations Research*, vol. 68, no. 1, pp. 89–108.

Saleth, R. M. and Dinar, A. 2004, *The institutional economics of water: a cross-country analysis of institutions and performance*, Edward Elgar Publishing.

Scheierling, S. M., Loomis, J. B. and Young, R. A. 2006, ‘Irrigation water demand: A meta-analysis of price elasticities’, *Water resources research*.

Scott, A. 2001, *The evolution of resource property rights*, Oxford University Press.

Seneviratne, S. I., Nicholls, N., Easterling, D., Goodess, C., Kanae, S., Kossin, J., Luo, Y., Marengo, J., McInnes, K., Rahimi, M. et al. 2012, ‘Changes in climate extremes and their impacts on the natural physical environment’. <http://ipcc-wg2.gov/SREX/>

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

Sharma, D. 2012, A New Institutional Economics Approach to Water Resource Management, PhD thesis, University of Sydney.

Shoham, Y., Powers, R. and Grenager, T. 2007, ‘If multi-agent learning is the answer, what is the question?’, *Artificial Intelligence*, vol. 171, no. 7, pp. 365–377.

Simon, H. A. 1965, ‘The architecture of complexity’, *General systems*, vol. 10, no. 1965, pp. 63–76.

Stachurski, J. 2008, ‘Continuous state dynamic programming via nonexpansive approximation’, *Computational Economics*, vol. 31, no. 2, pp. 141–160.

Stiglitz, J. E. 1975, ‘Incentives, risk, and information: notes towards a theory of hierarchy’, *The Bell Journal of Economics*, vol. pp. 552–579.

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

Takayama, T. and Judge, G. G. 1964, ‘Spatial equilibrium and quadratic programming’, *Journal of Farm Economics*, vol. 46, no. 1, pp. 67–93.

Tarlock, A. D. 2002, ‘The Future of Prior Appropriation in the West’, *Natural Resources Journal*, vol. 41, pp. 769.

Taylor, J. B. 1993, Discretion versus policy rules in practice, *in* ‘Carnegie-Rochester conference series on public policy’, Vol. 39, Elsevier, pp. 195–214.

Taylor, R. 2008, ‘Australian food bowl areas soon beyond repair’. <http://www.reuters.com/article/2008/06/18/us-australia-drought-idUSSYD3270620080618>

Tesauro, G. and Kephart, J. O. 2002, ‘Pricing in agent economies using multi-agent Q-learning’, *Autonomous Agents and Multi-Agent Systems*, vol. 5, no. 3, pp. 289–304.

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

The Wentworth Group of Concerned Scientists 2014, ‘National Water Reform’. <http://wentworthgroup.org/programs/national-water-reform/>

Timmer, S. and Riedmiller, M. 2007, Fitted q iteration with cmacs, *in* ‘Approximate Dynamic Programming and Reinforcement Learning, 2007. ADPRL 2007. IEEE International Symposium on’, IEEE, pp. 1–8.

Tisdell, J. 2010, ‘Acquiring Water for Environmental Use in Australia: An Analysis of Policy Options’, *Water Resources Management*, vol. 24, pp. 1515–1530.

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.

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.

Truong, C., Harris, M. and Drynan, R. 2010, A General Framework for Analysing Water Property Rights, *in* 'Australian Conference of Economists'.

United States Geological Survey 2009, 'Estimated Use of Water in the United States in 2005'.
<http://water.usgs.gov/watuse/>

URS Australia 2010, *Murray-Darling Basin Return Flows Investigation*.

Van Dijk, A., Evans, R., Hairsine, P., Khan, S., Nathan, R., Paydar, Z., Viney, N. and Zhang, L. 2006, *Risks to the shared water resources of the Murray-Darling Basin*.

Varian, H. R. 2014, 'Big data: New tricks for econometrics', *The Journal of Economic Perspectives*, vol. pp. 3–27.

Vaux, H. J. and Howitt, R. E. 1984, 'Managing water scarcity: An evaluation of interregional transfers', *Water Resources Research*, vol. 20, no. 7, pp. 785–792.

Verhoeven, H. 2012, 'Big is beautiful: Megadams, African water security, and China's role in the new global political economy'.
<http://www.globalwaterforum.org/2012/10/16/big-is-beautiful-megadams-african-water-security-and-chinas-role-in-the-new-global-political-economy/>

Victorian Ombudsman 2011, 'Investigation into the Foodbowl Modernisation project and other matters'.
<https://www.ombudsman.vic.gov.au/Publications/Parliamentary-Reports/Investigation-into-the-Foodbowl-Modernisation-Proj>

Watkins, C. and Dayan, P. 1992, 'Q-learning', *Machine learning*, vol. 8, no. 3, pp. 279–292.

Watson, A. 2005, Competition and water: a curmudgeonÂŠs view, *in* 'Australian Competition and Consumer Commission Conference,ÂŚRelationship Between Essential Facilities and Downstream MarketsÂŠ, Gold Coast, Queensland', Vol. 28.

Watson, A. 2007, ‘A national plan for water security: pluses and minuses’, *Connections: Farm, Food and Resources Issues.*

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

Weintraub, G., Benkard, C. and Van Roy, B. 2010, ‘Computational methods for oblivious equilibrium’, *Operations research*, vol. 58, no. 4-Part-2, pp. 1247–1265.

Weiring, M. and Otterlo, M., eds 2012, *Reinforcement learning: State-of-the-Art*, Vol. 12 of *Adaptaion, Learning and Optimization*, Springer.

Whiteson, S., Taylor, M. E., Stone, P. et al. 2007, *Adaptive tile coding for value function approximation*, Computer Science Department, University of Texas at Austin.

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.

Williamson, O. E. 1973, ‘Markets and hierarchies: some elementary considerations’, *The American Economic Review*, vol. pp. 316–325.

Willingham, R. 2011, ‘Significant financial hit to state, Baillieu warn’. <http://www.theage.com.au/victoria/significant-financial-hit-to-state-baillieu-warns-20110117-19u1f.html>

Wittfogel, K. 1957, *Oriental Despotism: A Comparative Study of Total Power*, Yale University Press.

Wolpert, D. H. 1996, ‘The lack of a priori distinctions between learning algorithms’, *Neural computation*, vol. 8, no. 7, pp. 1341–1390.

Wurbs, R. A. 1993, ‘Reservoir-system simulation and optimization models’, *Journal of water resources planning and management*, vol. 119, no. 4, pp. 455–472.

Wurbs, R. A., ed. 2013, *Water Resources Planning, Development and Management*, InTech.

Yeh, W. 1985, ‘Reservoir management and operations models: A state-of-the-art review’, *Water resources research*, vol. 21, no. 12, pp. 1797–1818.

Zetland, D. J. 2008, Conflict and cooperation within an organization: A case study of the Metropolitan Water District of Southern California, PhD thesis, University of California.

Appendix A

Parameterising the model

The parameter assumptions for the planner's version of the model (see section 3.5.3) are laid out below. Additional assumptions required for various decentralised versions of the model are documented in chapters 5, 6 and 7.

A.1 Data sources

A.1.1 MDB storages

Table A.2 contains statistics on 22 major storages in the MDB (those greater than 50GL in capacity — excluding the Snowy Mountains Scheme). This data set was compiled from a number of sources. Storage capacities, mean and median inflows were obtained from NWC (2011*a*). Additional variables (including surface area) were obtained from the ANCOLD (2013) register of large dams. Net evaporation data were matched to the dams using spatial climate data from the BOM (2013).

Annual inflow autocorrelation and coefficient of variation data were obtained from two sources: data for five major MDB rivers were obtained from (MDBA 2012*b*), data for a larger sample of Australian rivers were provided by Peel et al. (2010).

ABARES survey of irrigation farms

The ABARES survey of irrigation farms in the MDB began in 2006-07 (Ashton and Oliver 2012). The survey samples around 10 per cent of the irrigation farm population (around 900 farms) each year on a rotating basis (generating an unbalanced panel). Currently five years of data are available, 2006-07 to 2010-11, covering some extremely dry (2006-07 to 2008-9) and extremely wet (2010-11) conditions.

Hughes (2011) provides a detailed summary of the data set and describes methods for constructing key variables. All variables used in the regression analysis below are as defined in Hughes (2011), except for $INFLOW$ which is defined as $\tilde{I}_t$ from section 3.5.3. $INFLOW$ is calculated by catchment based on annual observed inflow.

Table A.1: $INFLOW$ values by catchment, 2006-07 to 2010-11

Year	Murrumbidgee	Goulburn/Loddon	Murray
2006-07	0.45	0.29	0.46
2007-08	0.44	0.60	0.42
2008-09	0.45	0.46	0.53
2009-10	0.66	0.82	0.78
2010-11	2.00	2.13	1.64

Table A.2: MDB storages greater than 50 GL

Name	River	State	Year constructed	Capacity (GL)	Median annual inflow (GL)	Mean annual inflow (GL)	Surface area ('000 Ha)	Mean annual net evap. (m)
Pindari Dam	Severn	NSW	1969	312	128	177	196	0.78
Copeton Dam	Gwydir	NSW	1976	1362	297	412	46	0.81
Keepit Dam	Namoi	NSW	1960	426	271	349	44	1.26
Split Rock Dam	Manilla	NSW	1987	397	31	70	22	1.10
Chaffey Dam	Peel	NSW	1979	62	41	53	5	0.66
Burrendong Dam	Macquarie	NSW	1967	1190	784	1009	89	0.94
Windamere Dam	Cudgegong	NSW	1984	368	40	55	20	0.75
Wyangala Dam	Lachlan	NSW	1971	1218	576	732	53	0.81
Burrinjuck Dam	Murrumbidgee	NSW	1928	1026	978	1254	55	0.55
Blowering Dam	Tumut	NSW	1968	1631	1531	1568	43	0.29
Hume	Murray	NSW	1936	3038	3889	4182	202	0.68
Menindee Lakes	Darling	NSW	1960	1731	904	1998	458	1.97
Dartmouth Dam	Mitta Mitta	VIC	1979	3908	850	892	63	0.09
Lake Eildon	Goulburn	VIC	1956	3334	1379	1467	138	0.24
Waranga Basin	Offstream	VIC	1905	432	1286	1485	58	0.93
Eppalock	Campaspe	VIC	1964	305.0	143	163	32	0.76
Cairn Curran	Loddon	VIC	1956	147.0	102	117	19	0.75
Tullaroop	Tullaroop Creek	VIC	1959	73.0	47	57	7	0.84
Glenlyon Dam	Pike Ck	QLD	1976	254.0	52	73	18	0.74
Leslie Dam	Sandy Ck	QLD	1965	106.0	22	31	13	0.89
Coolmunda	Macintyre Brook	QLD	1968	69.0	35	68	16	0.92
Beardmore	Balonne	QLD	1972	82.0	731	1161	29	1.32

A.2 Water supply

A.2.1 Storage capacity, K

As there are no scale effects in the model the absolute size of the river system is irrelevant. Storage capacity K is taken as the numeraire in the parameterisation and fixed at 1000 GL. Figure A.1 shows a histogram of capacities for the major MDB storages, the mean is 975 GL.

Figure A.1: Histogram of storage capacity, MDB storages

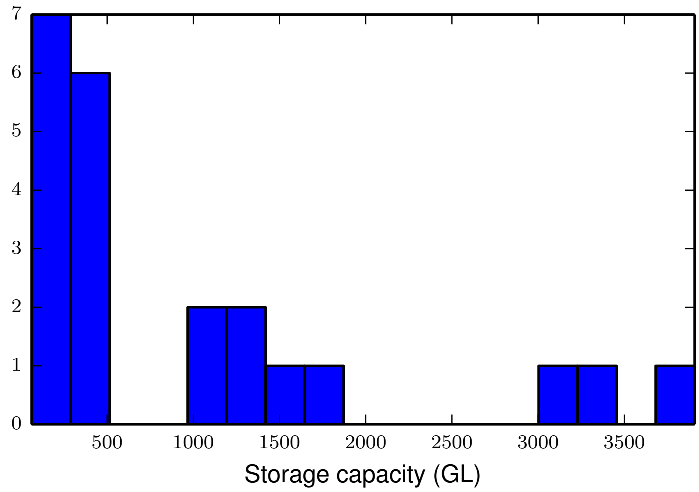

A.2.2 Mean inflow over storage capacity, $E[I_{t+1}]/K$

Figure A.2 shows a histogram of the ratio of mean annual inflow to capacity for the 22 MDB storages and our assumed distribution. Given K and ρ_I this ratio determines $E[\epsilon_t]$.

Figure A.2: Annual mean inflow over capacity, MDB storages

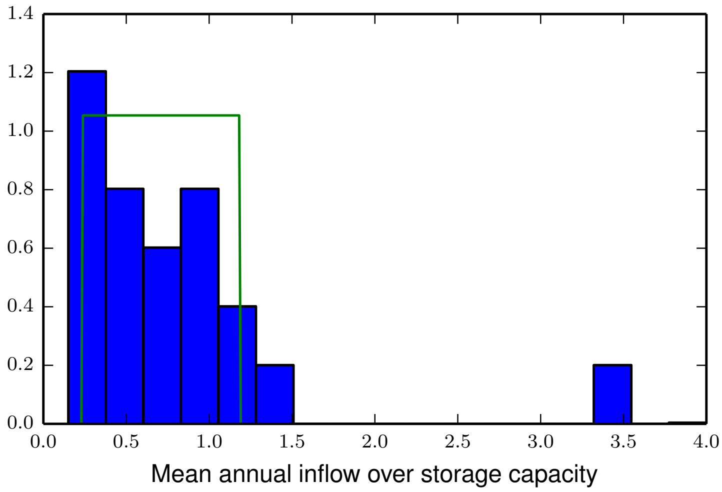

A.2.3 Coefficient of variation of inflow, c_v

Table A.3 shows the coefficient of variation c_v for annual inflows (the ratio of standard deviation to the mean) for five major MDB rivers (derived from modelled natural river flow time series MDBA 2012*b*). Figure A.3 shows a histogram of c_v for a larger sample of Australian rivers (obtained from Peel et al. 2010).

We assume a uniform distribution for c_v between 0.4 and 1.1. Given K, ρ_I, $E[\epsilon_t]$ and our assumption of a gamma inflow distribution, c_v determines k_I and θ_I.

Figure A.3: Standard deviation of annual inflow over mean, Australian rivers

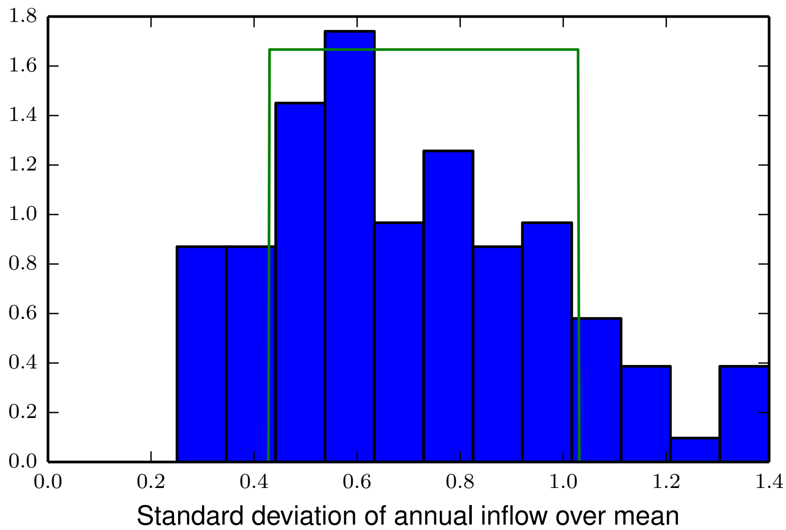

Table A.3: Annual inflow coefficient of variation, selected MDB rivers

River	Location	$\hat{\rho}_I$
Murray	Yarrawonga Weir	0.49
Murrumbidgee	Burrinjuck Dam	0.43
Goulburn	Lake Eildon	0.63
Namoi	Keepit Dam	0.88
Ballone	St George	1.08

A.2.4 Surface area over storage capacity, $\alpha K^{2/3}/K$

Figure A.4 shows the histogram for the ratio of surface area (HAs) to storage capacity (ML) for the 22 MDB storages and our assumed distribution. Given K this ratio determines the parameter α.

Figure A.4: Surface area over capacity, MDB storages

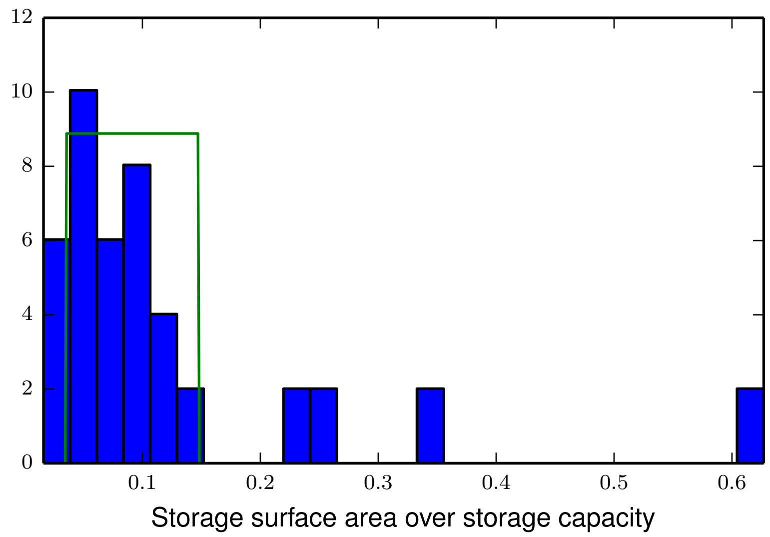

A.2.5 Net evaporation rate, δ_0

Figure A.5 shows a histogram for mean annual net evaporation[1] (meters) and our assumed distribution for the parameter δ_0.

A.2.6 Inflow autocorrelation, ρ_I

First order autocorrelation coefficients for five major MDB rivers are shown in table A.4 (derived from MDBA 2012*b*).

[1]Annual pan evaporation multiplied by a pan factor of 0.75 less mean annual rainfall.

Figure A.5: Mean annual net evaporation, MDB storages

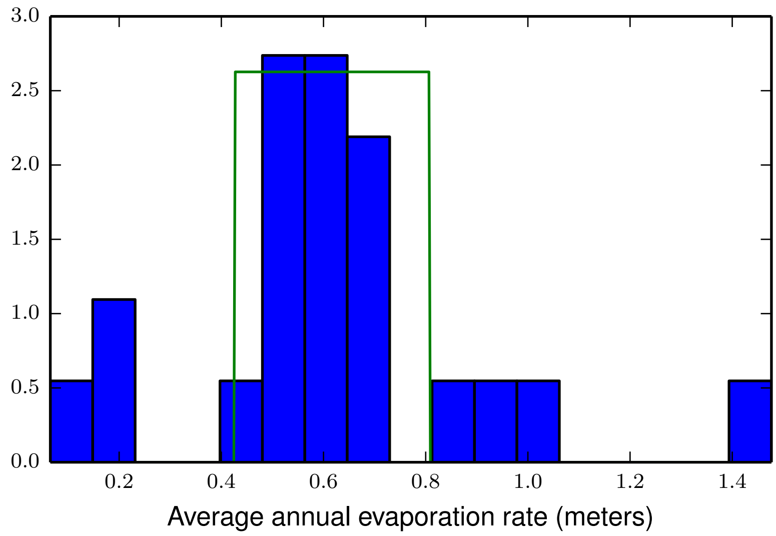

Table A.4: Estimated annual autocorrelation, selected MDB rivers

River	Location	$\hat{\rho}_I$
Murray	Yarrawonga Weir	0.23
Murrumbidgee	Burrinjuck Dam	0.28
Goulburn	Lake Eilden	0.28
Namoi	Keepit Dam	0.23
Ballone	St George	0.22

A uniform distribution for the parameter ρ_I is adopted in the range 0.2 to 0.3.

A.2.7 Delivery losses, δ_{1a}, δ_{1b}

Table A.5 shows estimates of fixed and variable irrigation system delivery losses for five major MDB irrigation areas. Estimates are derived by OLS from the data set of Hume (2008).

Table A.5: Estimated delivery loss coefficients, selected MDB irrigation areas

Irrigation Area	$\hat{\delta}_{1a}/E[I_t]$	$\hat{\delta}_{1b}$
Murray	0.05	0.15
Shepparton	0.04	0.25
Jemalong	0.04	0.25
Colleambally	0.01	0.23
Murrumbidgee	0.08	0.13

MDBA (2013), Gippel (2006) present estimates on river losses for the Murray River. Losses between Hume Dam and Yarrawonga are around 4 per cent on average with almost all losses independent of flow.

Uniform distributions are specified for $\delta_{0a}/E[I_t]$ on $(0.0, 0.15)$ and δ_{1b} on $(0.15, 0.30)$.

A.3 Water demand

A.3.1 Yield functions, θ_{h0} to θ_{h5}

Yield functions, mapping water use per unit land (ML / Ha) to profit per unit land ($ / Ha) were estimated using a sample from ABARES survey of irrigation farms in the MDB (Ashton and Oliver 2012) for the period 2006-07 and 2010-11. Central case estimates for $(\theta_{h0}, ..., \theta_{h6})$ are generated from the below regressions.

Broadacre farms

A broadacre revenue yield function was estimated for a sample of 378 southern MDB broadacre farms, using pooled OLS, see table A.6. A separate cost function was also estimated, see table A.7. Note the $MATERIALS$ variable excludes any water purchase costs.

Wine grape farms

A quantity yield function was estimated for a sample of 325 southern MDB wine grape areas (each farm observation can have multiple crop areas, see Hughes 2011). Results are shown in A.8. An estimated horticulture farm cost function is shown in table A.9.

Generating the θ parameters

In the high reliability case the quantity yield function is multiplied by the mean sample grape price ($597 per tonne) and the lagged yield effect is added to the current period yield effect (using a discount rate of 0.95). For both cases, profit function parameters are obtained by subtracting cost function coefficients from revenue function coefficients. All excluded variables ($CAPITAL/LAND$ etc.) were fixed at median sample values. Final estimates are shown in table A.12.

Figure A.6 displays the estimated yield function for low reliability users (based on southern MDB broadacre farms), figure A.7 shows the high reliability class (based on southern MDB wine grape areas). Final parameter estimates are shown in table

A.10. The parameters θ_{h0} to θ_{h5} are assumed to be normally distributed around these estimates, subject to some constraints. Standard deviations are derived from regression standard errors (see table A.12).

Table A.6: Southern MDB broadacre farm revenue yield function

Dependent variable: REVENUE/LAND			
Explanatory variables	Coefficient	Standard Error	P-value
Constant	88.84	47.86	0.06
$WATER/LAND$	398.95	51.11	0.00
$(WATER/LAND)^2$	-59.40	16.82	0.00
$CAPITAL/LAND$	0.28	0.04	0.00
$(CAPITAL/LAND)^2$	0.00	0.00	0.00
$INFLOW$	41.52	96.34	0.67
$INFLOW^2$	14.87	38.79	0.70
$INFLOW.(WATER/LAND)$	-48.48	17.32	0.01
$AREA1/LAND$	189.70	46.03	0.00
$AREA2/LAND$	63.56	8.26	0.00
$(CAPITAL/LAND)(WATER/LAND)$	-0.12	0.07	0.07
$(CAPITAL/LAND)(WATER/LAND)^2$	0.07	0.02	0.01
$(AREA2/LAND)(WATER/LAND)$	-20.48	9.28	0.03
$(AREA1/LAND)(WATER/LAND)$	-94.35	46.20	0.04
$MURRAY$	-20.85	23.54	0.38
$MURRUMBIDGEE$	-91.39	22.83	0.00
$GOULBURN$	54.50	28.67	0.06
R-squared	0.46		
Observations	707		
Cross sections	378		
Periods	5		
Residual autocorrelation	0.5		

Table A.7: Southern MDB broadacre farm cost function

Dependent Variable: MATERIALS/LAND			
Explanatory variables	Coefficient	Standard Error	Prob.
$Constant$	-41.97	24.83	0.09
$WATER/LAND$	101.63	15.39	0.00
$CAPITAL/LAND$	0.22	0.03	0.00
$(CAPITAL/LAND)^2$	0.00	0.00	0.03
$INFLOW$	20.82	10.83	0.06
$MURRAY$	37.01	18.40	0.04
$MURRUMBIDGEE$	33.30	19.17	0.08
$GOULBURN$	69.95	22.30	0.00
$AREA1/LAND$	207.31	29.17	0.00
$AREA2/LAND$	22.07	5.35	0.00
$AREA11/LAND$	4992.96	1317.46	0.00
$AREA12/LAND$	65.24	252.72	0.80
$AREA13/LAND$	168.05	57.98	0.00
$AREA14/LAND$	192.41	31.94	0.00
$AREA15/LAND$	182.94	30.01	0.00
R^2	0.47		
Observations	707		
Cross sections	378		
Periods	5		

Table A.8: Southern MDB wine grape area quantity yield function

Dependent Variable: (QTY8/AREA8)			
Explanatory variables	Coefficient	Standard Error	P-value
$Constant$	-1.12	2.22	0.61
$VOL8/AREA8$	3.61	0.51	0.00
$(VOL8/AREA8)^2$	-0.21	0.04	0.00
$INFLOW$	3.61	3.97	0.36
$INFLOW^2$	-0.90	1.60	0.57
$INFLOW(VOL8/AREA8)$	-0.20	0.25	0.43
$AREA8/LAND$	1.86	0.92	0.04
$LAND$	-4.9E-04	3.0E-04	0.10
$CAPITAL$	1.5E-06	6.1E-07	0.01
$CAPITAL/LAND$	-1.3E-04	6.6E-05	0.06
$VOL8(-1)/AREA8$	0.97	0.27	0.00
$(VOL8(-1)/AREA8)^2$	-0.05	0.02	0.00
R^2	0.49		
Observations	325		
Cross sections	165		
Periods	4		
Residual autocorrelation	0.36		

Table A.9: Southern MDB horticulture farm cost function

Dependent Variable: $MATERIALS/LAND$			
Explanatory variables	Coefficient	Std. Error	P-value
$Constant$	131.68	286.82	0.65
$WATER/LAND$	156.37	42.10	0.00
$CAPITAL/LAND$	0.23	0.04	0.00
$(CAPITAL/LAND)^2$	0.00	0.00	0.16
$INFLOW$	229.66	593.92	0.70
$INFLOW^2$	65.03	275.03	0.81
SA	335.89	216.66	0.12
$GOULBURN$	-887.13	338.70	0.01
$AREA4/LAND$	6127.73	503.05	0.00
$AREA5/LAND$	4218.32	492.93	0.00
$AREA6/LAND$	1424.30	463.62	0.00
$AREA7/LAND$	2810.88	575.95	0.00
$AREA8/LAND$	772.81	378.43	0.04
$AREA9/LAND$	8730.21	3373.52	0.01
R^2	0.40		
Observations	1065		
Cross sections	534		
Periods	5		

Figure A.6: Estimated yield function, low reliability users

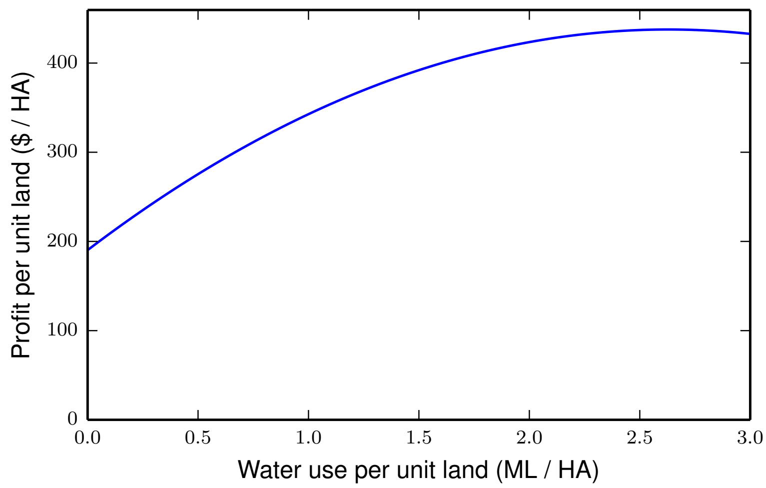

Figure A.7: Estimated yield function, high reliability users

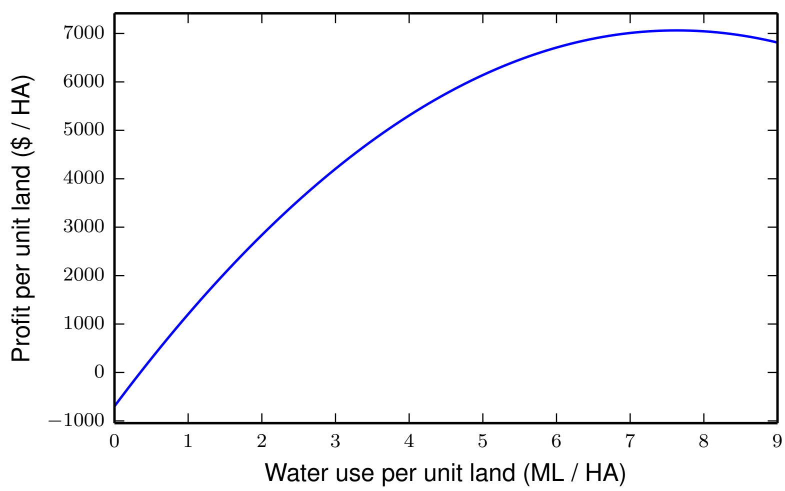

Table A.10: Estimated yield function parameters

Parameter	$h = low$	$h = high$
θ_{h0}	154.7	-1773.8
θ_{h1}	236.7	2135.0
θ_{h2}	-35.8	-133.3
θ_{h3}	20.7	1597.1
θ_{h4}	14.9	-520.9
θ_{h5}	-48.5	-100.8

A.3.2 Productivity shocks, ρ_e, σ_η

Productivity shocks are based on analysis of residuals from the estimated yield functions. Since the ABARES data set is an (unbalanced) panel autocorrelation within the residuals can be used to estimate ρ_e. A uniform distribution for ρ_e is specified over the range (0.3, 0.6).

A.3.3 Number of users, mix of high and low reliability users

In all parameterisations the number of users is fixed at 100. In the central case $n_{low} = n_{high} = 50$. A uniform distribution for n_{high} is specified over [30, 70].

Table A.11 shows horticultural water use as a proportion of total water use, for various MDB regions in 2011-12. The ratio of high to low farm land area $\mathscr{A}_{high}/\mathscr{A}_{low}$ is calibrated so that in the central case (where $n_{high} = 50$), high reliability users demand 20 per cent of the full supply volume (when $W_t = K$). This gives a value of for $\mathscr{A}_{high}/\mathscr{A}_{low}$ of 0.07 which, incidentally, is closely representative of the ratio of mean land areas of horticulture and broadacre farms in the ABARES data set.

Table A.11: Irrigation water use (GL), MDB and selected NRM regions, 2011-12

Region	Horticulture	All agriculture	Proportion
MDB	960.1	5,875.4	0.16
Southern MDB			
Murrumbdigee	119.5	1,256.8	0.10
Lower Murray	64.3	77.5	0.83
Goulburn-Broken	58.7	494.4	0.12
Lower Murray	64.3	77.5	0.83
Goulburn-Broken	58.7	494.4	0.12
North Central (VIC)	50.1	544.8	0.09
SA MDB	270.8	309.7	0.87
Northern MDB			
Border rivers (NSW)	4.9	436.8	0.01
Border rivers (QLD)	11.9	527.5	0.02
Namoi	0.7	253.3	0.00
Condamine	3.8	183.5	0.02

A.3.4 Aggregate demand / supply balance

The aggregate level of land relative to mean inflow (i.e., $\mathscr{A}_{low}$) is drawn from a uniform distribution with range $[5115.3, 8525.4] \times E[I_t]/K$. Under the central case value of $\mathscr{A}_{low}$, the expected water price in full supply $W_t = K$ year is \$10 per ML, reflective of water prices in the MDB during wet periods.

A.4 The complete parameterisation

A.4.1 Water supply

$$K = 1000000$$

$$\rho_I \sim U[0.20, 0.30]$$

$$E[I_t]/K \sim U[0.23, 1.18]$$

$$c_v = \frac{\sqrt{VAR[I_t]}}{E[I_t]} \sim U[0.40, 1.00]$$

$$E[\epsilon_t] = \frac{E[I_t]}{1 - \rho_I}$$

$$VAR[\epsilon_t] = VAR[I_t](1 - \rho_I^2)$$

$$k_I = \left(\frac{E[\epsilon_t]^2}{VAR[\epsilon_t]}\right)$$

$$\theta_I = \left(\frac{VAR[\epsilon_t]}{E[\epsilon_t]}\right)$$

$$\alpha.K^{2/3}/K \sim U[0.03, 0.15]$$

$$\delta_0 \sim U[0.43, 0.81]$$

$$\delta_{1a}/K \sim U[0.00, 0.15]$$

$$\delta_{1b} \sim U[0.15, 0.30]$$

A.4.2 Water demand

$$\theta_{h,k} \sim N[\mu_{\theta hk}, \sigma_{\theta hk}]$$

Subject to:

$$\theta_{h1} > 0$$

$$\theta_{h2}, \theta_{h3} < 0$$

Table A.12: Yield function parameters

Parameter	$h = low$		$h = high$	
	μ	σ	μ	σ
θ_{h0}	154.7	0.0	-1773.8	0.0
θ_{h1}	236.7	17.0	2135.0	102.1
θ_{h2}	-35.8	5.6	-133.3	7.4
θ_{h3}	20.7	0.0	1597.1	0.0
θ_{h4}	14.9	0.0	-520.9	0.0
θ_{h5}	-48.5	5.8	-100.8	50.5

$$0.50 \leq \bar{q}_{low} \leq 6.50$$

$$5.00 \leq \bar{q}_{high} \leq 14.00$$

$$\bar{q}_h = \frac{\theta_{h,1} + \theta_{h,3}}{-2\theta_{h,2}}$$

$$\rho_\epsilon \sim U[0.3, 0.5]$$

$$\sigma_\eta \sim U[0.1, 0.2]$$

Finally $\mathscr{A}_h$, n_h satisfy:

$$n_{high} + n_{low} = n = 100$$

$$n_{low} \sim U[30, 70]$$

$$\mathscr{A}_{low} / \mathscr{A}_{high} = 0.07$$

$$\frac{\mathscr{A}_{low}}{E[I_t]/K} \sim U[5115.3, 8525.4]$$

A.4.3 Discount rate

$$\beta \sim U[0.93, 0.96]$$

A.4.4 The central case

The central case parameterisation is defined as the mean values of all parameter distributions:

$$K = 1000000$$

$$k_I = 1.22$$

$$\theta_I = 434107.83$$

$$\rho_I = 0.25$$

$$\alpha = 9.08$$

$$\delta_0 = 0.62$$

$$\delta_{1a} = 53156.06$$

$$\delta_{1b} = 0.22$$

Table A.13: Central case θ parameters

Parameter	$h = low$	$h = high$
θ_{h0}	154.7	-1773.8
θ_{h1}	236.7	2135.0
θ_{h2}	-35.8	-133.3
θ_{h3}	20.7	1597.1
θ_{h4}	14.9	-520.9
θ_{h5}	-48.5	-100.8

$$\rho_\epsilon = 0.40$$

$$\sigma_\eta = 0.15$$

$$n = 100$$

$$n_{low} = 50$$

$$n_{high} = 50$$

$$\mathscr{A}_{low} = 4833.90$$

$$\mathscr{A}_{high} = 316.90$$

$$\beta = 0.945$$

Appendix B

Computation

B.1 Introduction

This appendix provides further detail on the implementation of the model and the algorithms used to solve it. Here we address some practical aspects of coding and computing the model, which were omitted from the main chapters.

We begin with an overview of the model code. We then provide some additional detail on the implementation of our model including: water account reconciliations, spot market clearing and optimal inflow share searches. Finally, we provide some more detail on the machine learning tools introduced in chapter 8, including tile coding and sample grids.

B.2 The code

This section is best read in conjunction with the model code, which is available online at:

[**github.com/nealbob/regrivermod**](github.com/nealbob/regrivermod)

B.2.1 Python and Cython

For this thesis we use a combination of Python and Cython. Python is a popular 'high-level' language, designed to provide simple easy to read code. Unfortunately, Python code runs slow compared with 'low-level' languages like C.

Cython is essentially a Python to C translator. Given 'Python like' code — roughly Python plus type declarations — Cython generates C code. Cython provides a language almost as simple as Python, but with performance near that of compiled C code. As a fast Object Orientated language, Cython is a viable alternative to C++ or Java, for this type of agent-based / multi-agent modelling.

Python code files have the extension `.py` and Cython the extension `.pyx`.

B.2.2 Overview

Most of the computationally intensive parts of the model are contained in two Cython based modules `econlearn` (a machine learning toolkit) and `regrivermod` (a simulation model of a regulated river).

Above these modules, all of the parameter assumptions are contained in `para.py`. For a given set of parameters `model.py` combines `econlearn` and `regrivermod` to solve the various versions of the model. The python scripts `chapter3.py, ...` implement sensitivity analysis. Finally, the `results` module is used to generate all of the figures and tables.

Figure B.1: Top level code structure

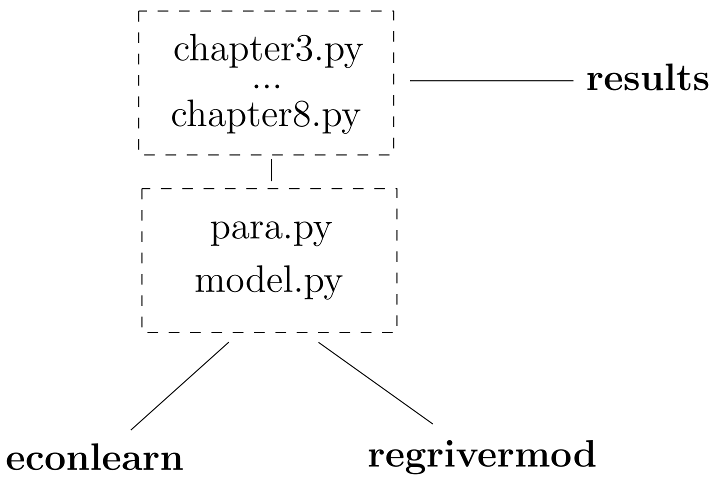

regrivermod

The main purpose of `regrivermod` is to perform Monte Carlo simulation and record data.

`regrivermod` contains the following classes:

- `Storage`: contains all of the hydrological detail of the model: inflows, storage, river flows and losses etc.

- `Utility`: plays the role of a water utility, implements water accounting / property rights
- `Users`: represents the consumptive water users, contains their policy, demand and payoff functions etc.
- `Market`: solves the spot market for the clearing price
- `Environment`: represents the environmental water holder, contains their policy, demand and payoff functions etc.
- `Simulation`: combines all of the above classes to perform simulations and record data
- `SDP`: a class for Stochastic Dynamic Programming, used to solve the planner's storage problem

econlearn

`econlearn` is a machine learning toolkit. Its main purposes is to implement the batch reinforcement learning algorithm: fitted Q-V iteration, using tile coding for function approximation. `econlearn` is designed to be a standalone tool, that could be used to solve other types of single / multi-agent economic problems.

`econlearn` contains the following files:

- `Qlearn.py`: implements fitted Q-V iteration. `Qlearn.py` accepts state / action samples as input and produces estimates of value and policy functions
- `tilecode.pyx`: implements tile coding function approximation
- `tile.py`: a python 'wrapper' for `tilecode.pyx`
- `samplegrid.pyx`: implements the distance based sample grid method

B.2.3 Parallelization

There are two forms of parallel computing: shared memory (i.e., threads) and message passing (i.e., processes).

With threading, multiple CPU cores can simultaneously operate on shared input data. With processes, each CPU core runs an essentially independent task, requiring its own copy of the input data. Message passing is suited to larger jobs — as there is overhead incurred in passing inputs and outputs. Threading is suited to smaller tasks such as, spreading a single loop over multiple cores.

We use both forms of parallelization in the model. Within Python, process based parallelization is achieved with the `multiprocessing` module. We use this module to run our model simulations in parallel. For example, with four CPU cores we can divide a 100,000 period simulation into four 25,000 period simulations, then combine the samples into a single data set (for more detail see **nealhughes.net/parallelcomp**).

Python doesn't allow for threading based parallelization, due to the 'Global Interpreter Lock' (GIL). However, multi-threading can be achieved with Cython via OpenMP (a threading platform for C). This type of parallelization is suited to loops where the order of computation does not matter. We use this type of multi-threading to speed up the fitting and prediction stages of tilecoding approximation (for more detail see **nealhughes.net/parallelcomp2**).

All of the central case results (and reported computing times) were generated using a Dell Optiplex 780, with a four core i7 CPU and 4 GB of RAM. The sensitivity analysis was completed with the aid of a small (10000 core hour) allocation on the National Computing Infrastructure (NCI) supercomputer. Each node on the supercomputer has 16 CPU cores, so we run four policy scenarios on a node at a time (with each using four cores). We then run a large number of such jobs, spread across multiple nodes (for more detail see **nealhughes.net/usingtheNCI**).

B.3 The regulated river model

B.3.1 Water accounting

The storage right accounting rules (introduced in chapter 5) are implemented in `regrivermod.utility.update_storage_accounts()`.

A small share of storage capacity is reserved to satisfy fixed delivery losses. Notionally this volume of capacity fl is held by the water utility. In chapter 5 $fl = \delta_{1a}/(1-\delta_{1b})$. In chapter 7 $fl = 2\delta_a + \delta_{Ea}/(1-\delta_{Eb})$.

The total volume available to water right holders is $K - fl$. In the rare event that $S_t < fl$ all account volumes s_{it} are set to zero. Any remaining water is effectively held in reserve to satisfy next periods fixed losses.

The final stage of the account update is an iterative procedure, which ensures the sum of water account volumes $\sum_{i=1}^{n} s_{it}$ equals the available storage volume

$S_{it} - fl$. Any discrepancy is shared across users according to inflow shares λ_i (i.e., socialised). Under capacity sharing scenarios this procedure computes any 'internal spills'. Internal spills need to be computed iteratively, because an initial round of internal spills, may lead further accounts to reach capacity creating further internal spills and so on.

The procedure is summarised in algorithm 9. In the actual implementation x_{it} are capped once user accounts reach their maximum (or minimum). This is not essential but it speeds up the convergence and makes x_{it} more interpretable.

Algorithm 9: Final water account reconciliation

initialise *tol*
set $x_{it} = 0$ for all i
set $\Delta = (S_t - fl) - \sum_{i=1}^{n} s_{it}$
while $|\Delta| > tol$ **do**
    set $x_{it} = x_{it} + \Delta\lambda_i$
    set $s_{i,t+1} = \max\{\min\{s_{it} - w_{it} - l_{it} + \lambda_i I_t + x_{it}, k_t\}, 0\}$ for all i
    $\Delta = (S_t - fl) - \sum_{i=1}^{n} s_{it}$
end

B.3.2 The spot market clearing price

With a transaction cost, it's not possible to derive the spot market clearing price analytically. Further, for any given level of aggregate supply Q_t and inflow I_t the clearing price can vary depending on the prevailing distribution of e_{it} and a_{it}.

Solving for the clearing price is then a numerical route finding problem in one dimension. We begin with the *secant method*. Occasionally this method fails to converge: typically when the clearing price is near the kink (between high and low reliability demand) in the market demand curve. In this case we turn to the slower but more reliable *bisection method*.

The code for this is contained in `regrivermod.market.solve_price()`. The secant method is summarised in algorithm 10.

To obtain good starting values (P_1) we first solve for a large number of (Q_t, I_t) points and fit a tile coding approximation to the market demand function.

B.3.3 The inflow share search

The inflow share search (see chapter 6) is a one dimensional unimodal optimisation problem. For this we use a simple stochastic hill climbing method. For the release sharing scenarios the method works similar to algorithm 11.

Algorithm 10: Secant method for obtaining a market clearing price

```
initialise $Q_t, tol, EX_0, EX_1 EX_2, P_1, P_2, P_{max}$
while $|EX_0| > tol$ do
    set $P_0 = \max\{\min\{P_1 - \frac{EX_1}{(EX_0 - EX_1)}(P_1 - P_2), P_{max}\}, 0\}$
    set $EX_0 = \sum_{i=1}^{n} d_i^{-1}(P_0, I_t, e_{it}) - Q_t$
    set $P_2 = P_1$
    set $P_1 = P_0$
    set $EX_2 = EX_1$
    set $EX_1 = EX_0$
end
```

Algorithm 11: Stochastic hill climbing inflow share search

```
initialise $\Lambda_{high}, \Lambda^*_{high}, SW, SW_{max}$
set $\delta = \Lambda_{high}/1.5$
set $ITER = 12$
for $i = 0$ to $ITER$ do
    draw $\epsilon \sim U(0,1)$
    set $\Lambda_{high} = \Lambda^*_{high} + \epsilon\delta$
    solve model and record social welfare $SW$
    if $SW > SW_{max}$ then
        set $\delta = 0.8\delta$
        set $\Lambda^*_{high} = \Lambda_{high}$
        set $SW_{max} = SW$
    else
        set $\delta = -0.8\delta$
    end
end
```

In the decentralised case we use a similar method, only we don't completely re-solve the model each time, we just run the learning algorithm for a fixed number of iterations (8): giving the users time to adapt to the new inflow share levels. In the case where inflow and storage capacity are 'unbundled' we have a two dimensional problem. Here we use a *move one at a time* version of the method.

All the code for these searches is located in `model.chapter6(.)`.

B.4 Machine learning methods

B.4.1 Tile coding

Computing tile indexes

With a single layer of evenly spaced tiles, the tile weights can be stored in a multi-dimensional array. For example, with D input dimensions (regressors) indexed by $k \in \{0, 1, ..., D\}$ and T_k tiles per input, our weight array would have shape $(T_0, ..., T_D)$. We then compute indexes by just scaling the input data to the range $(0, T_k)$ on each dimension and applying integer truncation.

With multiple layers our array would have shape $(N_L, T_0, ..., T_D)$ where N_L is the number of layers. To compute indexes we then need to include each layer's displacement or offset term. Function `index()` below computes array indexes, given an input point $\mathbf{X} = (X_0, ..., X_D)$, a layer index $i \in (0, 1, ..., N_L)$ and a displacement vector $\mathbf{dv}$ of shape (N_L, D).

Function index($\mathbf{X}$, i)

```
for k = 0 to D do
    set X_k = T_k(X_k − a_k)/(b_k − a_k) ;     // scale X_k from (a_k, b_k) to (0, T_k)
    set j_k = int(X_k + dv_ik)
end
return j
```

In practice, we stack the weights from each dimension and each layer of the scheme into a single dimensional array. The `index()` function in `econlearn.tilecode.pyx` computes the indexes to this array.

Hashing

Using the above method requires a weight array of size $\bar{N} = T_0 \times T_1 \times ... \times T_D \times N_L$. One way of lowering the memory requirements is to store the weights in a 'hash table' (i.e., an array of key value pairs).

The idea is that a large number (in some cases the vast majority) of tile weights may be inactive (i.e., zero) — as discussed in chapter 8 inputs tend to be correlated leaving large parts of the input space empty. In this case, we can significantly reduce the memory requirements by storing only the non-zero weights in a hash table. This is *perfect hashing* in the sense that no data is lost.

Our tile coding scheme implements a custom hash table using *open addressing* and *linear probing*. With hashing we need some *hash function* to map the indexes of the large array (lets call it A) to indexes of the smaller array (B). In some cases, the hash function will map the two A indexes onto the same B index: a *collision*. When writing to the array, any collisions are resolved by searching through array B one step at a time, until we find an empty slot. When reading we follow the same procedure, searching until we find the record with the matching key (the index of A).

We use the following hash function $h : [0, 1, ..., \bar{N}] \rightarrow [0, 1, ..., \underline{N}]$ for $\bar{N} > \underline{N}$:

$$h(x) = x(x + 3) \mod \underline{N}$$

Hash indexing is implemented in the `hash_idx()` function in `econlearn.tilecode.pyx`.

This light weight hashing scheme involves only a small loss of speed for a large reduction in memory use. Given the number of dimensions and tiles, most of our tile coding functions don't require hashing. However, it becomes essential when computing sample grids by tile coding (section B.4.2) where we use a higher number of tiles per dimension.

Displacement vectors

We adopt the 'optimal' displacement scheme of Brown and Harris (1994). This scheme is designed to distribute the displacements evenly across the input space; in contrast with the standard uniform displacements which move along a diagonal (figure B.2).

Under the Brown and Harris (1994) scheme the displacement array $\mathbf{dv}$ is generated from an initial vector $\mathbf{d} = (d_0, ..., d_D)$. Tables provided in Brown and Harris (1994), list all of the arrays $\mathbf{d}$, for combinations of N_L and D.

Figure B.2: Example displacement vectors in two dimensions

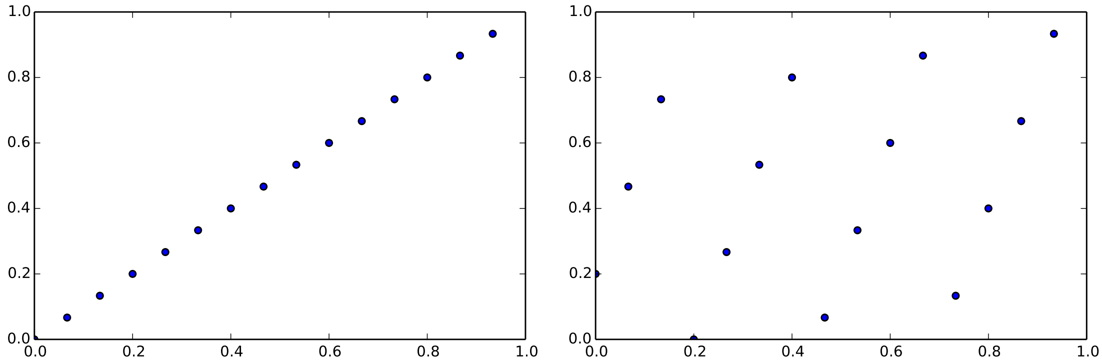

(a) uniform (b) ‘optimal’

Given **d**, **dv** is computed as follows

Algorithm 12: Computing optimal displacement vectors

for $i = 0$ *to* N_L **do**
  for $k = 0$ *to* D **do**
    set $dv_{ik} = 1 - N_L^{-1}((i+1)d_k \mod N_L)$
  end
end

Linear spline extrapolation

When using tile coding to approximate policy and value functions we also estimate a linear spline model for extrapolation. This model is sparse in the sense that there are no interaction terms

$$\hat{Y} = \sum_{k=1}^{D} \sum_{j=1}^{TL} \beta_{kj} \max\{X_k - TL^{-1} \times j, 0)\}$$

Here TL is the number of (evenly spaced) spline *knots* and TL^{-1} is the distance between them. Note here $\mathbf{X}$ is first scaled to $(0, 1)$. In practice, we typically set $TL = 4$ and $T_k = 10$. The spline model is fit by OLS.

B.4.2 Sample grids

The distance based sample grid method introduced in section 8.3.4 is effective for small samples (i.e., single agent problems), but becomes time consuming for large samples (i.e., multi-agent problems).

The first option is to add an early stopping condition (to algorithm 4): stop searching if N consecutive points fail to yield any new grid points (all are within distance $\underline{r}$ of an existing grid point).

A much larger speed-up can be obtained by employing tile coding approximation, since we can avoid calculating distances. Here we define a tile code (i.e., overlapping grid) *data structure* with tile widths of $\underline{r}$.

The algorithm works similar to before, only instead of calculating distances between the new data point and all existing grid points, we make one call to a tile code scheme. If any of the tiles overlapping the new point are 'active' then the new point is discarded, if not it is added to the gird.

As shown in figure B.3, for each grid point the active tiles approximate a hypercube of width $2\underline{r}$ around that point. Only points outside the cube are added to the grid.

Figure B.3: Tile code data structure for sample grids

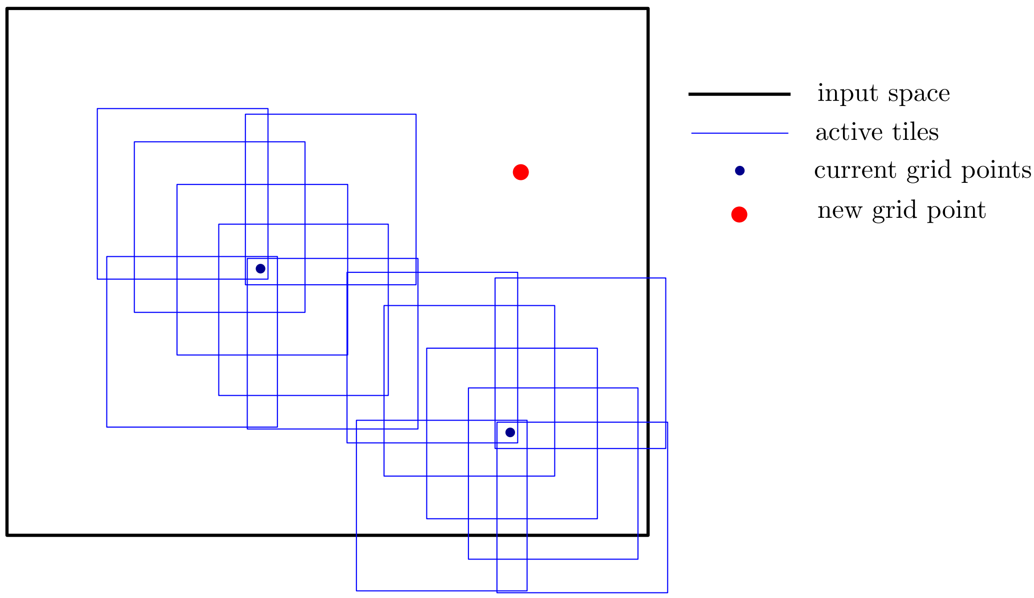

B.4.3 Random Forests

The field of supervised learning involves a massive variety of techniques all suited to different problem types. While there is no first-best method, random forests is perhaps the most popular technique: because it is known to perform well in most problem types with minimal tuning.

Random forests is an *ensemble* method: a random forest prediction is an average of a large number of *decision trees*. Each decision tree produces a piecewise constant approximation: with each segment defined as the average of the points contained within it. A random forest is then closely related to tile coding (when fit by averaging), only instead of a grid data structure there are *trees* (i.e., nested if statements). The advantage of trees over grids is that they can handle large dimensional spaces.

In random forests a large number of decision trees are estimated with randomised *splits*. There are a variety of approaches used to construct these trees. In this thesis we employ the method of *extremely randomized trees* (Geurts et al. 2006), as implemented in the python module `scikit-learn` (Pedregosa et al. 2011).

Appendix C

More on storage rights

C.1 Storage rights in the MDB and western US

C.1.1 Murray-Darling Basin

Currently, all MDB regulated rivers have some form of user level storage right. Approaches to storage rights can be classified into four systems: 'carry-over rights' (southern NSW), 'spillable water accounts' (northern VIC), 'continuous accounting' (northern NSW) and 'capacity sharing' (southern QLD), see table C.1 for a summary or Hughes et al. (2013) for a very detailed discussion.

A key problem in the southern MDB is the annual accounting for water withdrawals. Since user accounts are only updated for withdrawals annually, user account balances do not reflect actual storage levels. As such it difficult to internalise storage capacity constraints with account limits. In response, southern NSW and northern VIC impose arbitrary rules, such as limits on annual carryover volumes (account balances transferred between water accounting years). However, these rules are clumsy and often fail to bind, resulting in close to open access outcomes (Hughes et al. 2013).

A prime example, is the Victorian Murray region during the spill events of 2010-11 and 2011-12. For various reasons (see Hughes et al. 2013) spill forfeit rules were not applied to user accounts. This effectively allowed open access storage despite binding capacity constraints. This then resulted in substantial externalities, as users rushed to exploit the situation by trading unused water from NSW into Victorian water accounts (Hughes et al. 2013).

The 'continuous accounting' systems of northern NSW involve more frequent water withdrawal accounting, such that storage constraints can 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 very 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 (an 'announced allocation' system), which can lead to policy uncertainty. Finally, the northern Victorian approach is not precisely represented by any single scenario as it combines aspects of the SWA, CS-SWA and OA scenarios.

Table C.1: Storage and inflow rights in the MDB

	Carryover rights	Spillable Water Accounts	Continuous Accounting	Capacity Sharing
Region	southern NSW	northern VIC	northern NSW	southern QLD
Withdrawal accounting period	Annual	Annual	Monthly	Daily
Announced allocation / Inflow rights	Allocation	Allocation	Allocation	Inflow rights
Flow rights	Priority	Priority	Priority	Proportional
Annual carryover limits	Yes	Yes	No	No
Spill forfeit rules	No	No	Yes	No
Account limits	Yes	Yes	Yes	Yes
Account limit accuracy	Low	Low	Medium	High
Loss deductions	No	Yes	No	Yes

Source: Hughes et al. (2013)

C.1.2 Western US

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 reminiscent of the southern MDB (see chapter 4). Storage rights in the Texas Lower Rio Grande are similar to the continuous accounting systems of northern NSW, the approach in Southe Platte basin is 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 a combination of capacity shares and 'spill forfeit rules'. As an off-river reservoir, inflows into San Louis are pumped, so there are no physical spills. However, water is forfeited by users holding more than their capacity share to those with less than their share when the dam is full, in much the same way as northern Victoria.

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 (see Hughes 2013) they effectively allow irrigation districts to store unused water allocations in the Hoover Dam. Since the introduction of ICS over 1200 GL of water — around 10 per cent of the current balance of Hoover Dam — has accumulated under these rights (Hughes 2013).

C.2 Sample simulation results

Figures C.1 to C.4 below show sample simulation results for the CS, SWA, OA and NS scenarios, under the central case parameters. Each figure shows a sample of 50 periods with identical shocks (so that the results are comparable across scenarios). These samples are generated after the completion of the learning algorithm (i.e., after users have had time to learn near optimal policies).

The top left panel shows aggregate storage S_t and withdrawals W_t, while the top right depicts user storage accounts s_{it} and withdrawals w_{it} (all in percentages). Here the high reliability users are grouped on the right (users 50 to 100) and low on the left. The width of the bars indicates the size of user shares λ_i — with high reliability users having smaller shares. Under CS and NS, user accounts have a maximum of 100 per cent. Under SWA and OA, individual accounts can exceed 100 per cent.

On the bottom left we have the market water price, and on the bottom right the user net trade positions, in percentages ($\frac{a_{it}-q_{it}}{\lambda_i K}$). Negative values indicate water purchases and positive values water sales. For a complete set of results see **nealhughes.net**.

A number of insights can be gained from these results:

- In general high reliability users are more conservative (maintain higher storage account balances)
- In general net trade is from low to high users during dry periods and vice versa
- Higher storage levels (and less volatile water prices) are observed under OA and SWA
- Lower storage levels (and more volatile water prices) are observed under NS
- NS allows for little heterogeneity in user withdrawals and therefore greater reliance on trade
- More user level variation in storage and withdrawals is observed under OA and SWA, relative to CS

Figure C.1: Sample simulation results for the CS scenario

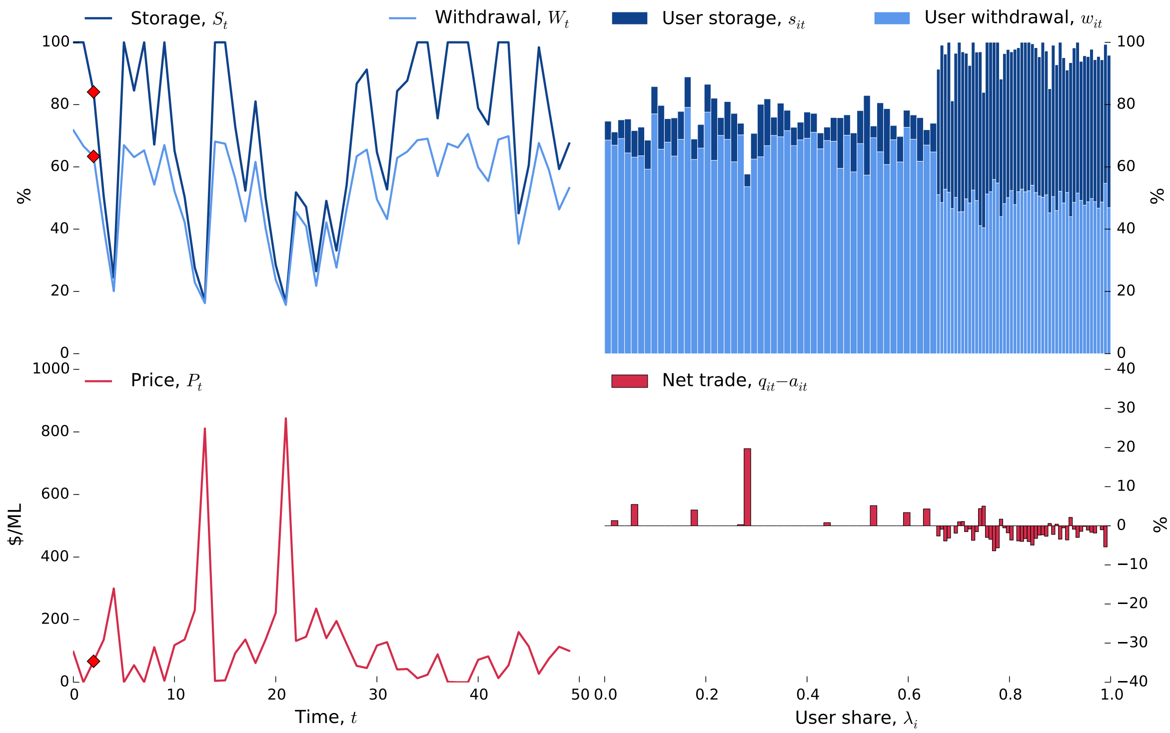

(a) $t = 2$

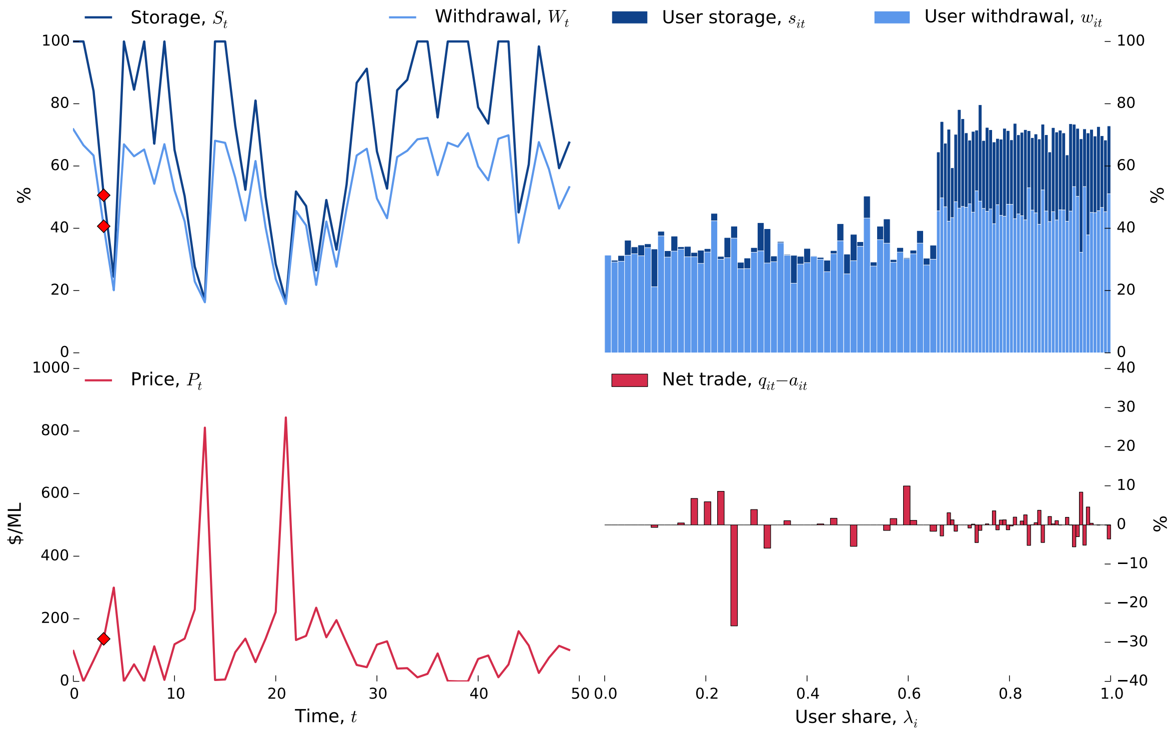

(b) $t = 3$

Figure C.1: Sample simulation results for the CS scenario

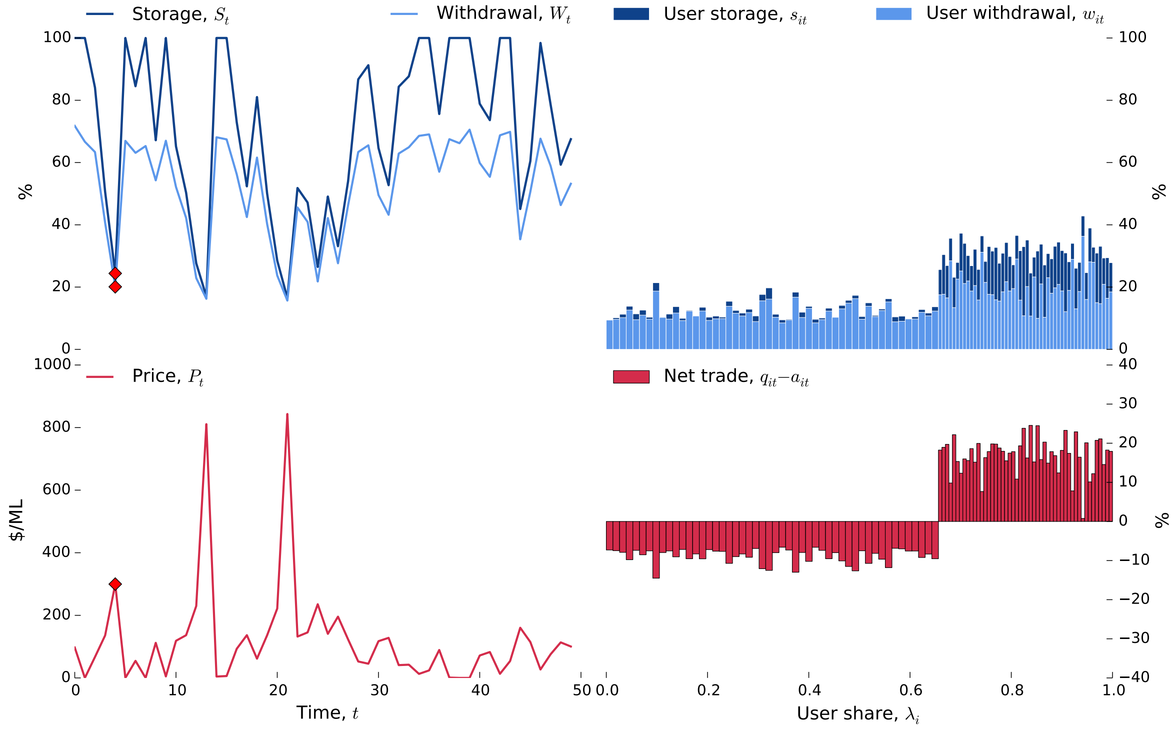

(c) $t = 4$

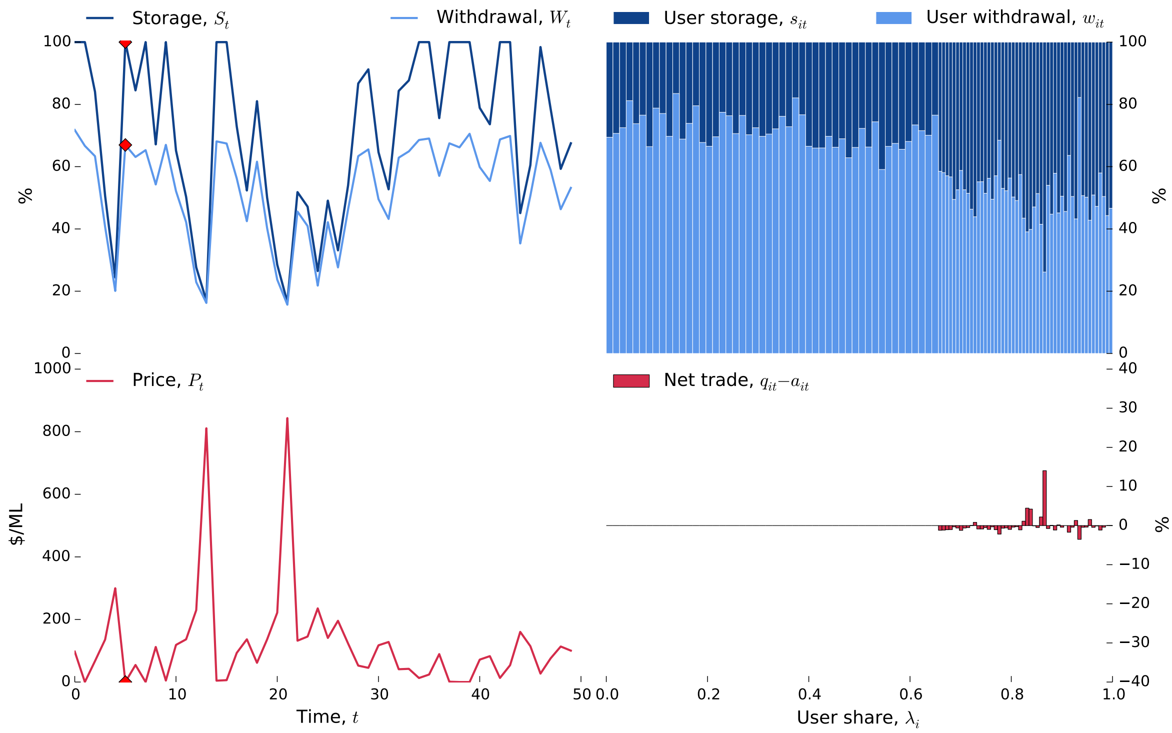

(d) $t = 5$

Figure C.2: Sample simulation results for the SWA scenario

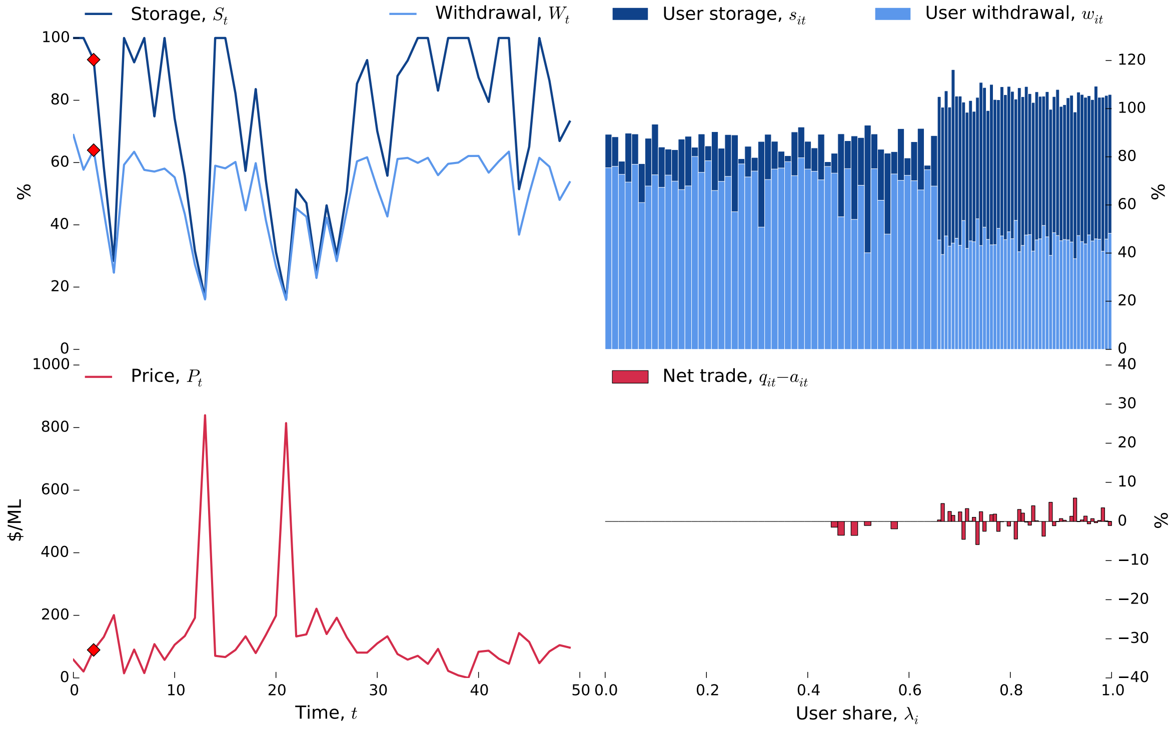

(a) $t = 2$

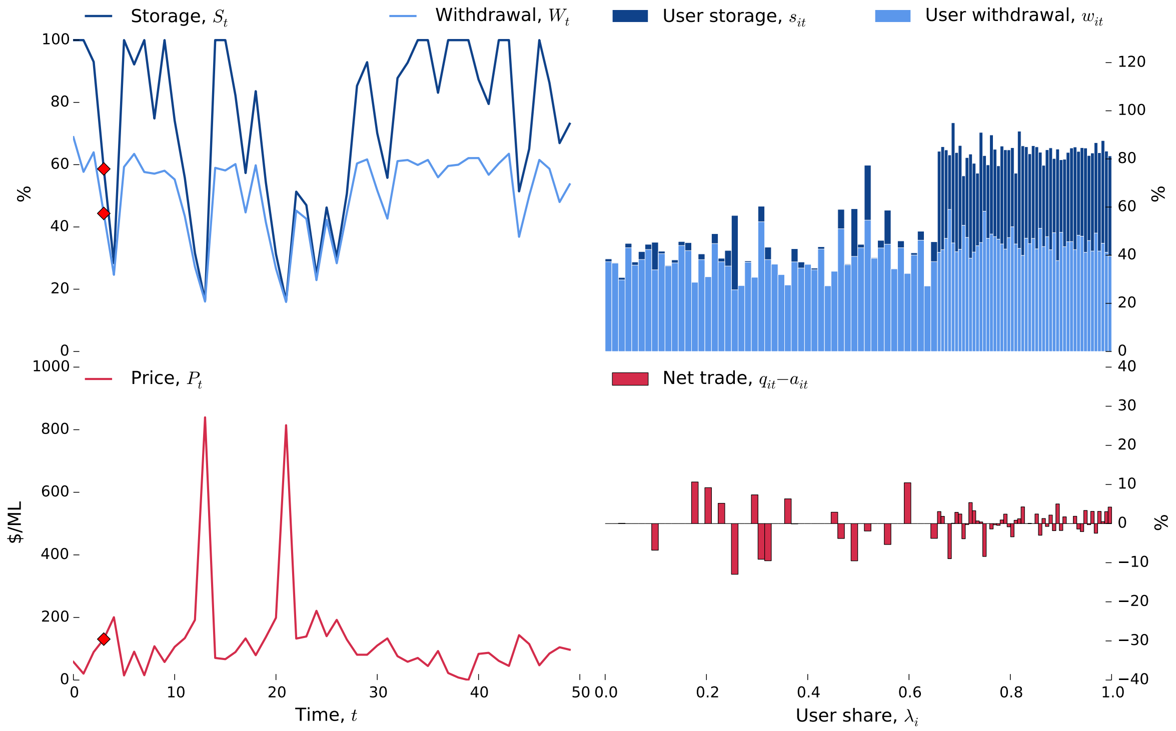

(b) $t = 3$

Figure C.2: Sample simulation results for the SWA scenario

(c) $t = 4$

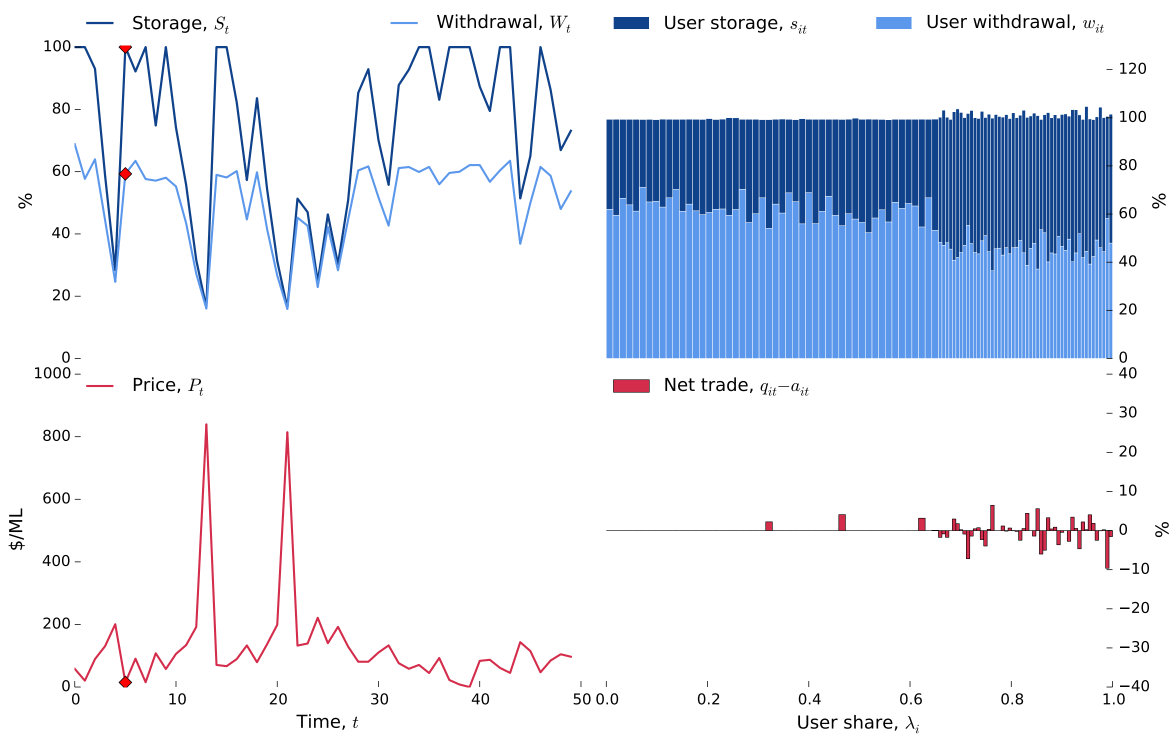

(d) $t = 5$

Figure C.3: Sample simulation results for the OA scenario

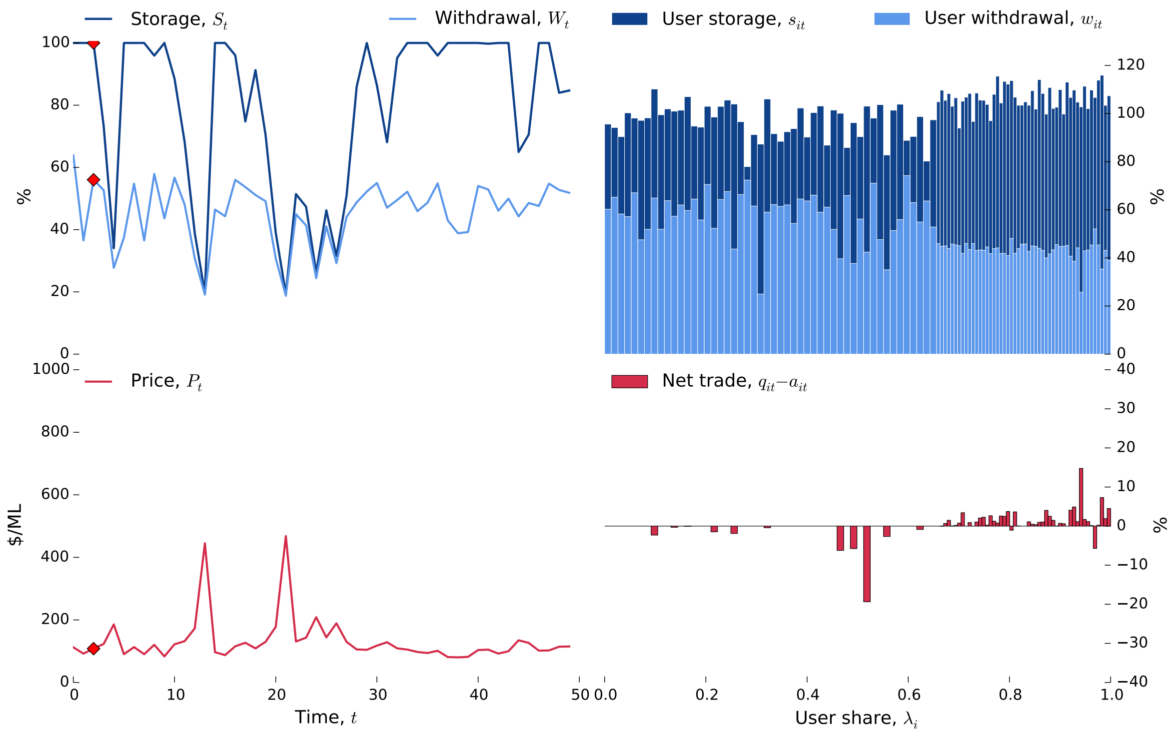

(a) $t = 2$

(b) $t = 3$

Figure C.3: Sample simulation results for the OA scenario

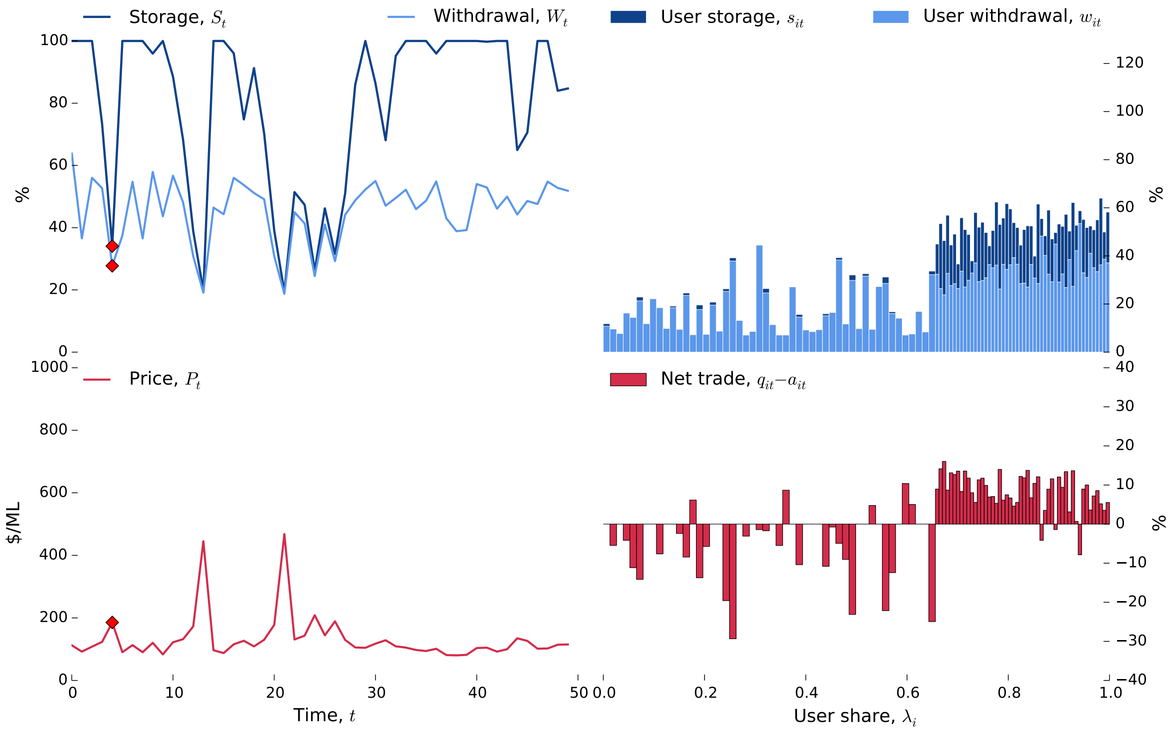

(c) $t = 4$

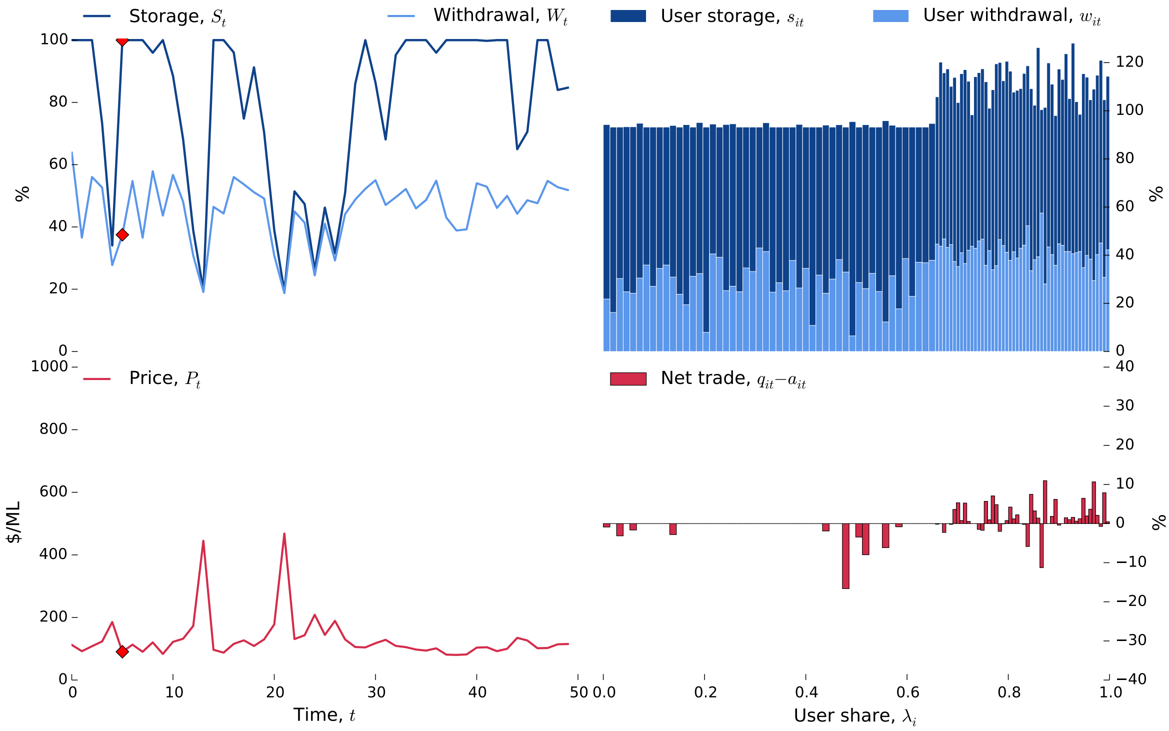

(d) $t = 5$

Figure C.4: Sample simulation results for the NS scenario

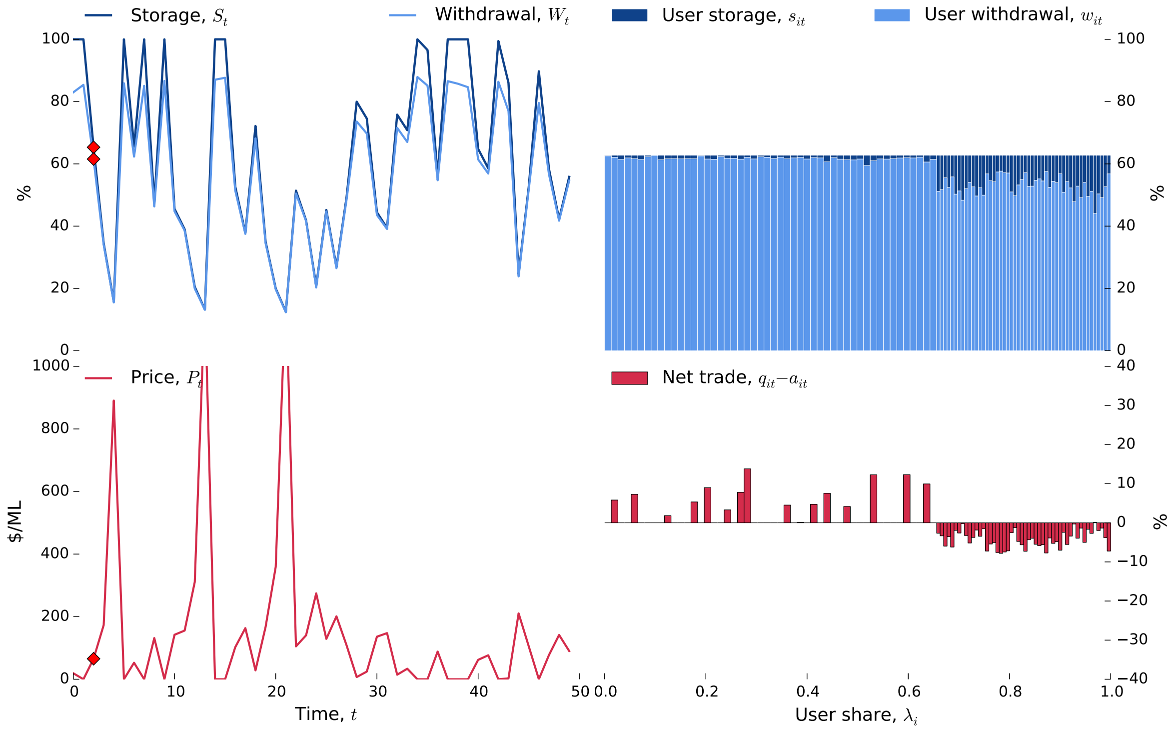

(a) $t = 2$

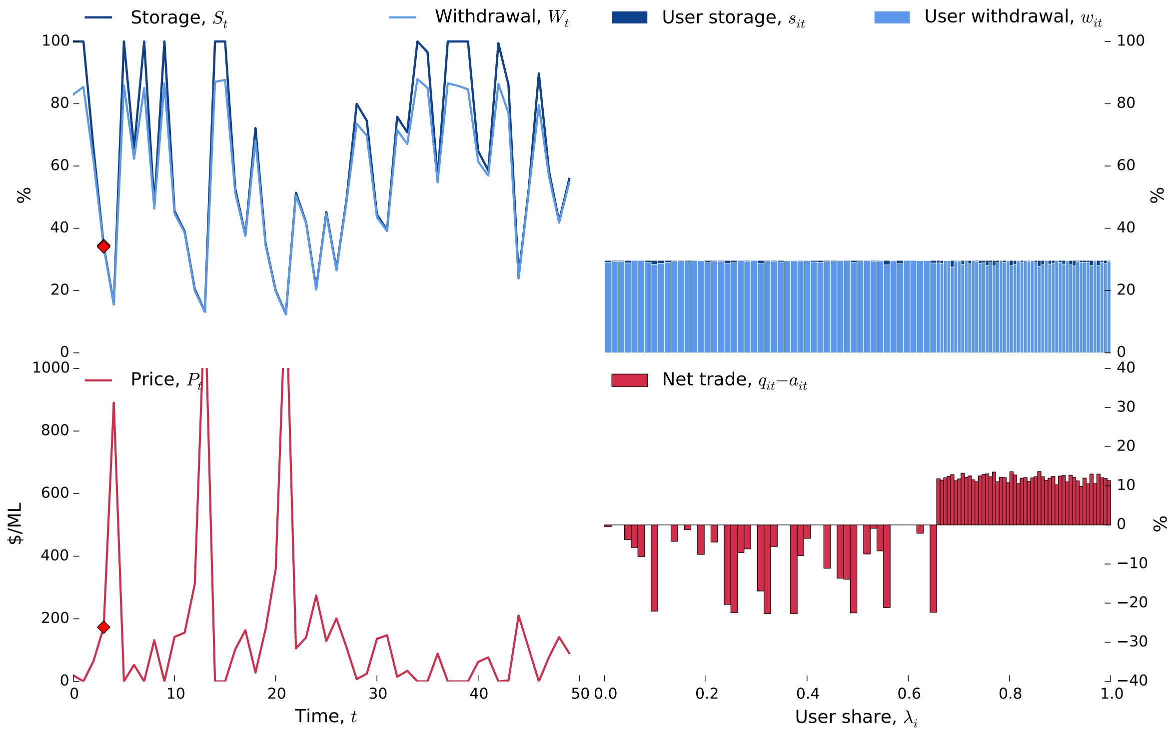

(b) $t = 3$

Figure C.4: Sample simulation results for the NS scenario

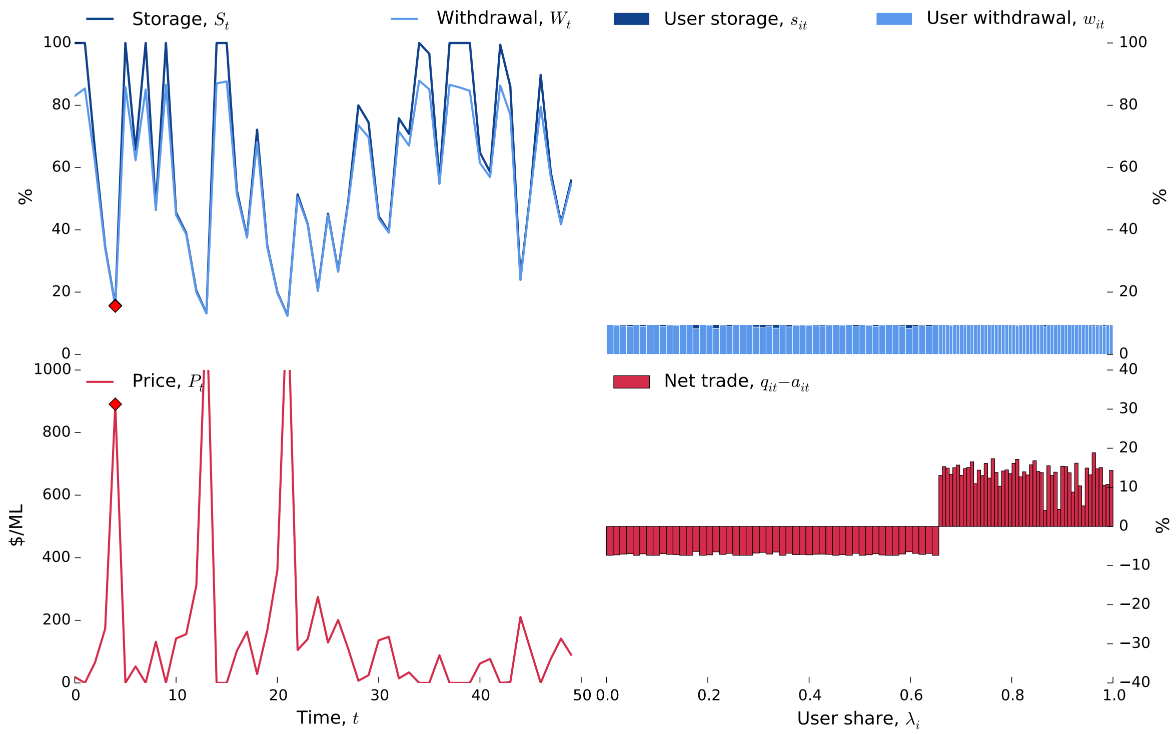

(c) $t = 4$

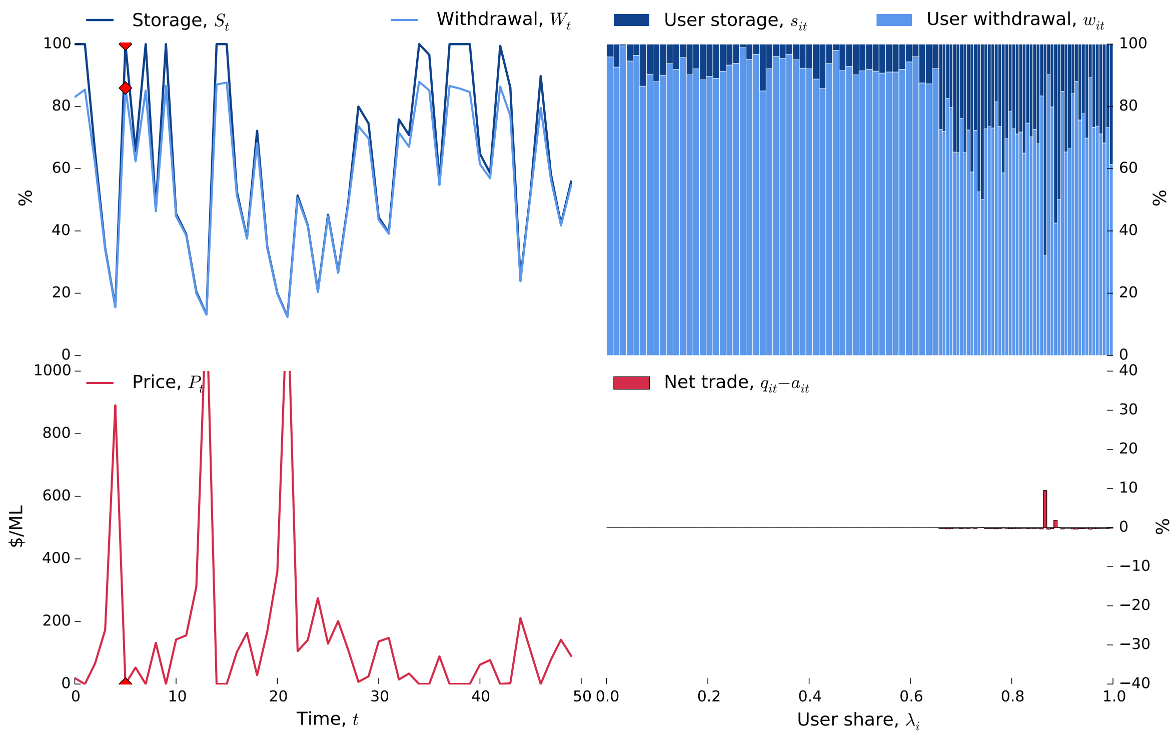

(d) $t = 5$

C.3 No trade scenario

Here we solve the model of chapter 5 with the central case parameters, but with no spot market (with $\tau = \infty$). The results are shown in tables C.2 to C.7.

Comparing these results with those of section 5.6.1 we can estimate the mean *gains from trade* (figure C.5). Here we see that in the presence of storage rights the gains from spot market trade are fairly small (in the order of \$2m a year) — because storage rights help to mitigate trade requirements (see chapter 6). In contrast, the gains from trade are large in the NS scenario.

In terms of storage, we see that the removal of trade actually reduces storage distortions. Generally, the scenarios achieve mean storage closer to the optimal level (see section 5.6.1) — for example OA involves slightly less over-storage and NS slightly less under-storage. This result is consistent with the argument of Brennan (2008*a*) that the introduction of trade can have a negative effect on inter-temporal efficiency by exacerbating storage externalities. However, in our case this negative effect is quite small such that trade has a net positive impact overall.

Figure C.5: Gains from trade, (\$m)

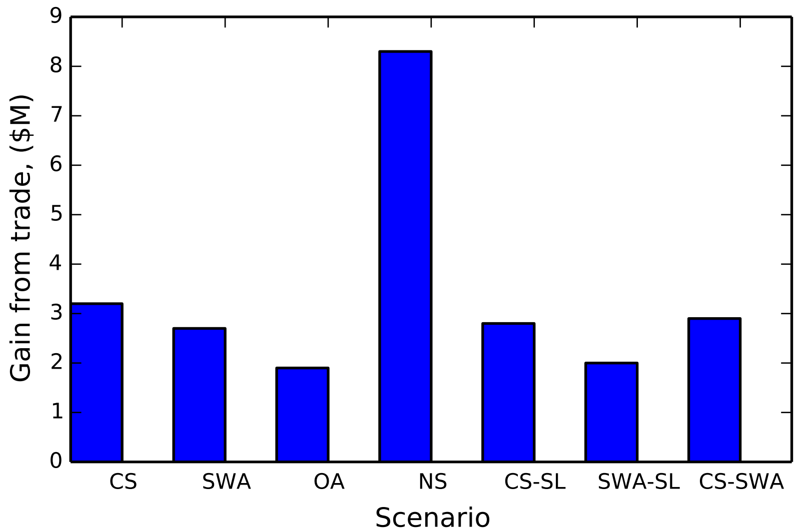

Table C.2: Social welfare, $\sum_{i=1}^{n} u_{it}$ (\$m)

	Mean	SD	2.5th	25th	75th	97.5th
CS	182.1	32.7	74.0	176.9	200.9	209.0
SWA	182.6	30.3	79.1	180.5	198.6	206.6
OA	180.8	26.0	93.6	178.8	194.1	202.5
NS	173.7	45.0	43.6	160.3	202.7	210.8
CS-SL	182.7	30.9	82.1	176.7	200.5	208.6
SWA-SL	183.2	27.8	91.8	179.7	198.1	205.9
CS-SWA	182.5	30.9	76.5	180.8	198.9	207.4
Planner	186.6	25.7	128.8	179.1	202.2	209.4

Table C.3: Storage, S_t (GL)

	Mean	SD	2.5th	25th	75th	97.5th
CS	698.3	282.7	143.2	466.9	1,000.0	1,000.0
SWA	728.9	277.9	150.3	509.3	1,000.0	1,000.0
OA	773.2	269.8	164.0	579.7	1,000.0	1,000.0
NS	611.6	299.6	105.8	353.8	942.3	1,000.0
CS-SL	706.2	279.4	148.4	480.2	1,000.0	1,000.0
SWA-SL	738.3	272.1	161.9	527.6	1,000.0	1,000.0
CS-SWA	725.0	280.5	146.1	501.8	1,000.0	1,000.0
Planner	697.9	282.3	160.4	462.2	1,000.0	1,000.0

Table C.4: Withdrawal, W_t (GL)

	Mean	SD	2.5th	25th	75th	97.5th
CS	520.4	170.5	139.4	398.4	675.4	696.4
SWA	499.0	139.6	146.3	421.6	604.3	642.8
OA	450.8	106.1	159.4	406.1	522.9	582.7
NS	562.0	252.7	102.4	351.4	824.2	896.3
CS-SL	518.2	166.8	145.9	398.5	666.8	710.3
SWA-SL	495.1	136.4	159.7	415.8	595.7	634.3
CS-SWA	500.8	141.6	142.7	423.5	597.5	659.2
Planner	521.0	176.0	158.4	390.3	666.5	690.5

Table C.5: Spills, Z_t (GL)

	Mean	SD	2.5th	25th	75th	97.5th
CS	130.4	279.8	0.0	0.0	88.7	1,066.3
SWA	150.3	300.3	0.0	0.0	151.6	1,122.3
OA	196.3	342.2	0.0	0.0	270.1	1,244.3
NS	92.7	232.8	0.0	0.0	0.0	972.6
CS-SL	133.2	282.1	0.0	0.0	99.2	1,078.5
SWA-SL	153.0	302.7	0.0	0.0	158.9	1,133.7
CS-SWA	149.8	299.8	0.0	0.0	150.2	1,116.2
Planner	129.9	278.5	0.0	0.0	88.8	1,051.6

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

	Mean	SD	2.5th	25th	75th	97.5th
CS	80.5	15.2	45.5	70.2	92.4	97.0
SWA	79.7	13.6	45.7	72.5	89.6	94.2
OA	77.3	11.1	45.8	75.4	84.0	90.8
NS	80.5	16.5	44.5	68.3	94.5	99.5
CS-SL	80.3	15.0	45.3	70.6	92.2	96.8
SWA-SL	79.3	13.8	44.8	71.8	89.4	93.9
CS-SWA	80.0	13.7	45.6	72.9	89.6	94.8
Planner	79.3	18.6	40.5	68.3	92.8	96.9

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

	Mean	SD	2.5th	25th	75th	97.5th
CS	101.7	21.0	25.3	102.2	110.9	115.3
SWA	102.8	19.9	30.3	103.5	111.2	115.5
OA	103.5	17.2	44.2	103.4	110.5	114.8
NS	93.2	30.1	-1.2	91.9	111.0	115.7
CS-SL	102.4	19.4	33.4	102.4	110.8	115.2
SWA-SL	103.9	17.7	42.1	103.9	111.1	115.3
CS-SWA	102.6	20.3	27.8	103.2	111.2	115.4
Planner	107.4	13.5	87.9	106.7	112.2	116.0

C.4 Earlier models

Early during this thesis a number of small models were developed for testing. While, these models are greatly simplified they are able to replicate some of the results from chapter 5 for the main storage right scenarios: CS, SWA and OA.

For completeness, we briefly summarise these models and their results below. Note that the assumptions, parameters and solution methods adopted here all differ significantly from our main model. The code used to generate these results is written in Matlab and can be provided on request.

C.4.1 Two time period model, no trade

This model is essentially that of chapter 5 but with only two time periods $t \in (0, 1)$, no spot market (user payoff functions π_i depend only on w_{it}), no evaporation or delivery losses and no inflow autocorrelation.

With two time periods user i's optimal strategy in period 1 is to consume all remaining water:

$$w_{i1}^* = \min\{s_{i0} - w_{i0} + \lambda_i I_1 + x_{i1}, k_{i1}\}$$

Now the optimisation problem for user i is

$$\max_{w_{i0}} \left(\pi_i(w_{i0}) + \beta E[\pi_i(\min\{s_{i0} - w_{i0} + \lambda_i I_1 + x_{i1}, k_{i1}\})]\right)$$

Subject to

$$0 \leq w_{i0} \leq s_{i0}$$

So now we have a static game with n players. The solution to user $i's$ problem is then a best response function h_i for w_{i0}

$$w_{i0}^* = h_i(., \mathbf{w}_{-i,0})$$

We solve for each user's best response function by simple grid search. We then iterate over the best response functions to find the (in this case unique) Nash equilibrium.

Below we present some numerical results for this problem with the following parameter assumptions:

- $K = 2$
- $\beta = 0.95$
- $I_1 \sim N_0^{4.5}(0.5, 1)$
- $\pi_i = W_{it}^{\alpha}.\mathscr{A}_i^{1-\alpha}$ where $\mathscr{A}_i$ is a fixed input and $\mathscr{A}_i = 1/n$ for all i.
- $\alpha_i = 0.2$ for half of the users and 0.6 for the other half
- $\lambda_i = 1/n$
- $S_0 = K$ and $s_{i0} = K/n$

Figure C.6 shows the best response functions for $n = 2$, for each of CS, SWA and OA.

Figure C.6: Best response functions for $n = 2$

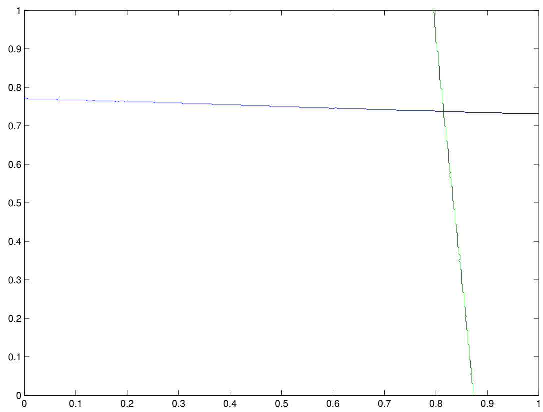

(a) SWA

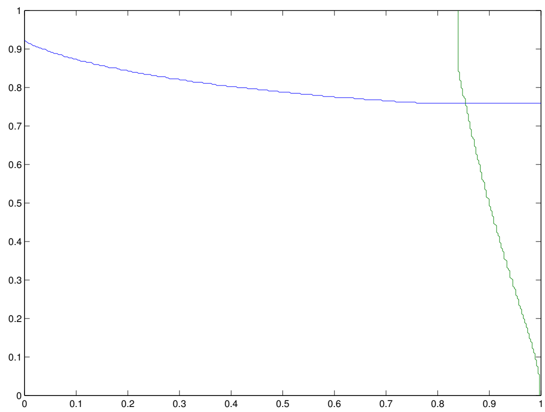

(b) CS

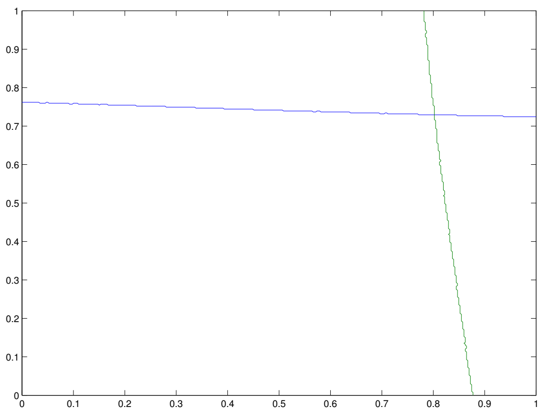

(c) OA

Table C.8 shows the equilibrium results for aggregate volume stored (i.e. $\bar{S}_0 - \bar{W}_0$) and aggregate welfare. Mean water storage is highest under OA, followed by SWA then CS. Welfare is highest under CS, followed by SWA and then OA. The differences in welfare and storage levels increase as the number of users increases.

Note that for the CS scenarios we assume 'uniform opponent behaviour', such that for each user i internal spills x_{i1} are defined

$$x_{i1} = \max\{\sum_{-i} \lambda_j I_1 - (\sum_{-i} \lambda_j K - \sum_{-i} s_{j0} + \sum_{-i} w_{j0}), 0\}$$

Given this assumption the CS results are independent of the number of users n.

Table C.8: Equilibrium results

	$n = 2$	$n = 4$	$n = 6$	$n = 50$
CS				
Storage	0.389	0.389	0.389	0.389
Welfare	2.252	2.252	2.252	2.252
SWA				
Storage	0.448	0.482	0.492	0.512
Welfare	2.250	2.249	2.249	2.248
OA				
Storage	0.469	0.515	0.535	0.569
Welfare	2.249	2.247	2.246	2.244

C.4.2 Two users infinite time horizon, no trade

In this version we have only two users $i \in (1, 2)$ but assume $t \in (0, 1, ..., \infty)$. Again we have no spot market, no evaporation or delivery losses and no inflow autocorrelation.

We solve this model for MPE using the standard value iteration algorithm, with a continuous approximation (using linear interpolation) to the value function. We solve for each of the CS, SWA and OA scenarios. A version of the social planner's problem with no water trading (RS — release sharing with proportional shares) is also solved.

The following parameter assumptions were made:

- $\beta = 0.95$
- $K = 2, \lambda_i = 0.5$
- $I_{t+1} \sim N_0^2(0.5, 1)$
- $\pi_i = a_i + b_i W_i + c_i W_i^2$
- $a_1 = 0, a_2 = -0.5, b_1 = 1, b_2 = 2.5, c_1 = -0.3, c_2 = -1.3$

Figure C.7: Payoff functions π_1, π_2

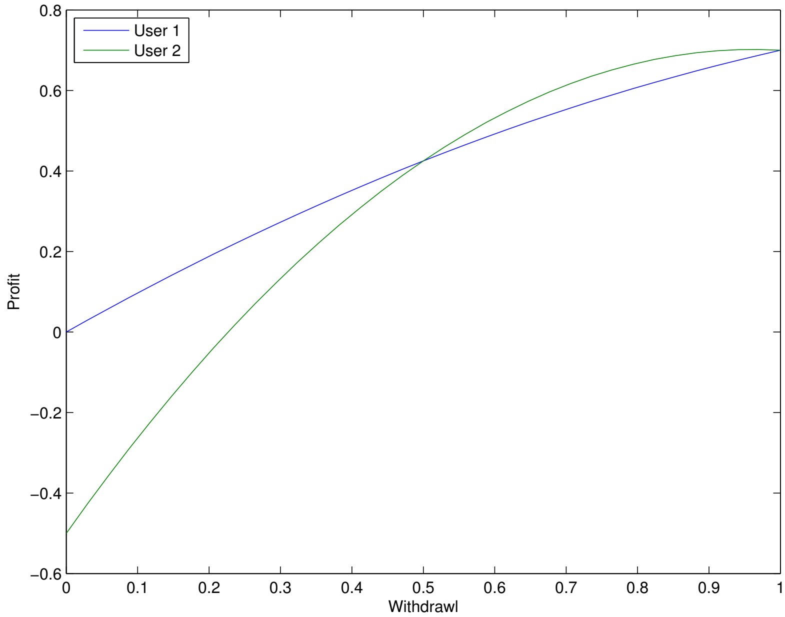

Estimated policy functions are shown in figures C.8 and C.9. User policies (withdrawals w_i) depend both on the users' state variable s_i and on the opponent state variable s_j.

Table C.9: Long run average results, storage levels and payoffs

	$\sum_i u_{it}$	u_1	u_2	S_t	s_{1t}	s_{2t}
RS	0.628	0.354	0.274	1.090	-	-
CS	0.643	0.358	0.284	1.047	0.492	0.556
SWA	0.644	0.357	0.286	1.09	0.505	0.586
OA	0.640	0.352	0.287	1.100	0.496	0.605

Resulting long run average payoff's and storage levels are summarised in table C.9, the results are broadly consistent with the two period model above.

Figure C.8: Social planner policy function — RS

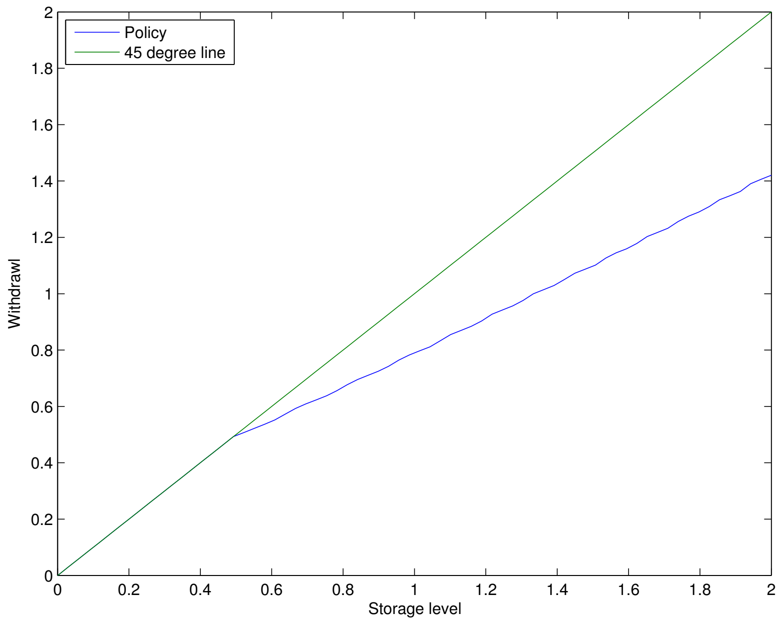

Figure C.9: User policy functions

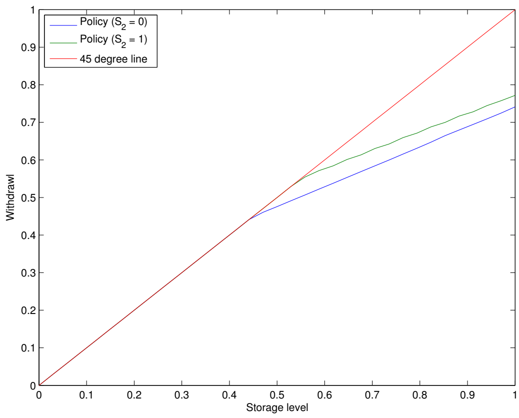

(a) CS, $i = 1$

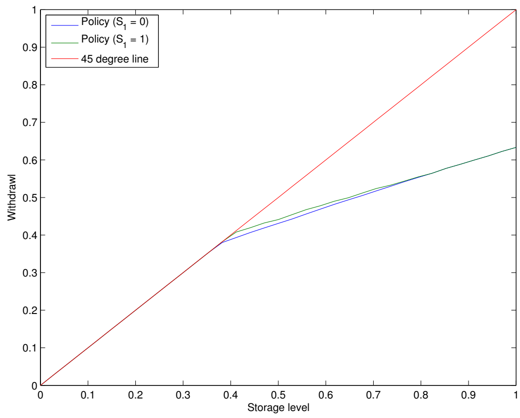

(b) CS, $i = 2$

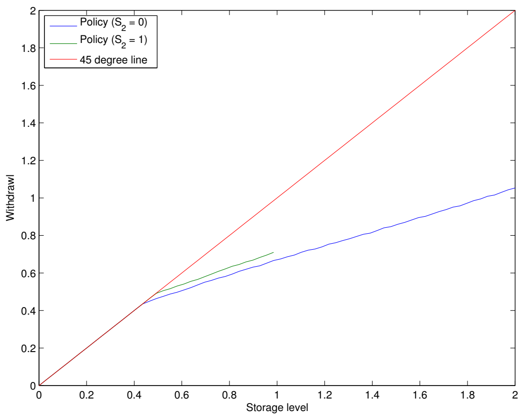

(c) SWA, $i = 1$

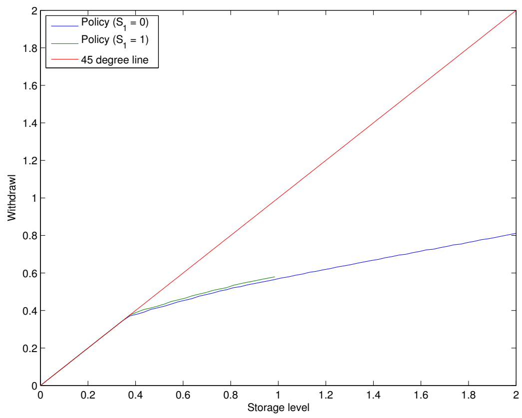

(d) SWA, $i = 2$

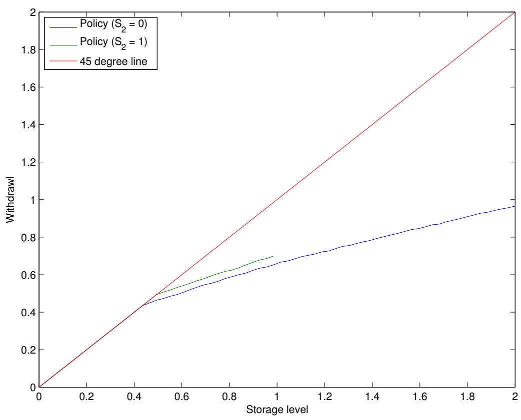

(e) OA, $i = 1$

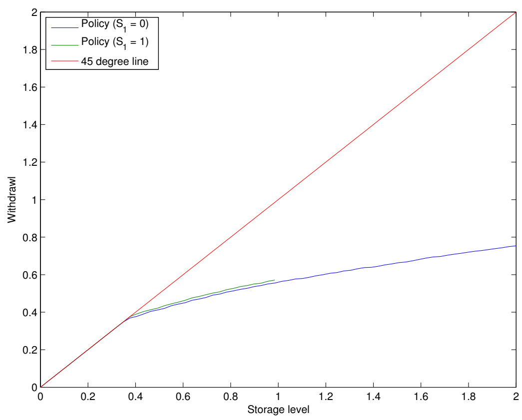

(f) OA, $i = 2$

Appendix D

More on priority rights

D.1 Introduction

This appendix presents some additional results on priority rights, for background refer to chapter 6.

D.2 Share conversion

D.2.1 The standard approach to priority rights

The standard approach to the two priority right system, involves a SOP type storage release rule:

$$W_t = \min\{S_t, \bar{S}\}$$

$$\bar{A} = \max\{\bar{S}(1 - \delta_{1b}) - \delta_{1a}, 0\}$$

where $\bar{S} < K$ is the maximum release volume and $\bar{A}$ the maximum allocation.

Now with a_{it} defined as in our RS-HL scenario, we can define the *nominal entitlement volume* (the maximum possible allocation) for each right holder $\bar{a}_i$ as

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

The *yield* of a water right is then defined as the mean per cent allocation

$$\frac{1}{T}\sum_{t=0}^{T}\frac{a_{it}}{\bar{a}_i}$$

Statistics like yield or reliability (the percentage of years in which a 100 percent allocation is received) are poor measures of the relative value of water rights. A better measure of value is the discounted value of future allocations[1].

$$\frac{1}{\bar{a}_i}\sum_{t=0}^{T}\beta^t a_{it} P_t$$

In this environment the planner has two policy parameters, Λ_{high} which determines the mix of rights and $\bar{S}$ which determines the storage policy. Or equivalently the volume of high reliability rights $\bar{A}\Lambda_{high}$ and the volume of low reliability rights $\bar{A}(1-\Lambda_{high})$.

A simple grid search of $\Lambda_{high}, \bar{S}$ reveals the optimal combination (for the central case parameters) as 0.163, 609 GL. This generates welfare of $185.0m compared the RS-HL-O scenario of $185.8m (where storage policy is fully flexible).

In this case, high reliability rights have a yield 0.95 and a value of $1,747 per ML, and low a yield 0.77 and value of $893 per ML.

D.2.2 Share conversion market

In a share conversion market different priority rights can be traded (either between users or via a central agency) at some exchange rate[2]. Share conversion, relies on the assumption that for any combination of $\Lambda_{high}, \bar{S}$ there exists an exchange rate between low and high rights, which preserves the characteristics (i.e., yield and / or value) of existing rights. If this is not the case the conversion of rights will lead to externalities on other right holders.

The idea is that a unit of high reliability is worth more than a unit of low. So as we convert low into high priority (i.e., increase Λ_{high}) we need an offsetting decrease in the total volume of water rights (i.e., a decrease in $\bar{S}$). In theory, the decrease in $\bar{S}$ ensures the additional high priority rights can be supplied without lowering the yield / value of existing rights.

If successful, such a market would allow both the mix of rights and the storage policy to be determined by market forces. However, there are some problems with

[1] We would expect this fundamental value to equate with the equilibrium price for permanent water rights in a zero transaction cost, risk neutral, perfect information environment.

[2] This exchange rate needn't be fixed but may vary depending on the prevailing mix of rights.

the approach. Firstly, even if it exists, this exchange rate will be difficult for a planner to calculate.

At best, a planner might hope to find the rate at which yield or reliability of rights is held constant. Of course this will not ensure that the value of rights is constant.

Even targeting yield would be difficult. In practice there will always be some form of storage right (even if it's the option to refuse delivery of water as in our NS scenario). As such, the yield of water rights is only partially under the control of planners.

Putting these problems aside, let's see if we can find the correct exchange rate under the SOP / RS-HL setup detailed above. Using our grid search results, figures D.1 and D.2 show the combinations of $\Lambda_{high}, \bar{S}$ that preserve constant yield of low and high priority rights respectively.

Clearly these two sets of lines do not overlap: there is no way to vary the mix of rights while maintaining the yield of both classes constant.

Figure D.1: Constant yield combinations, high priority rights

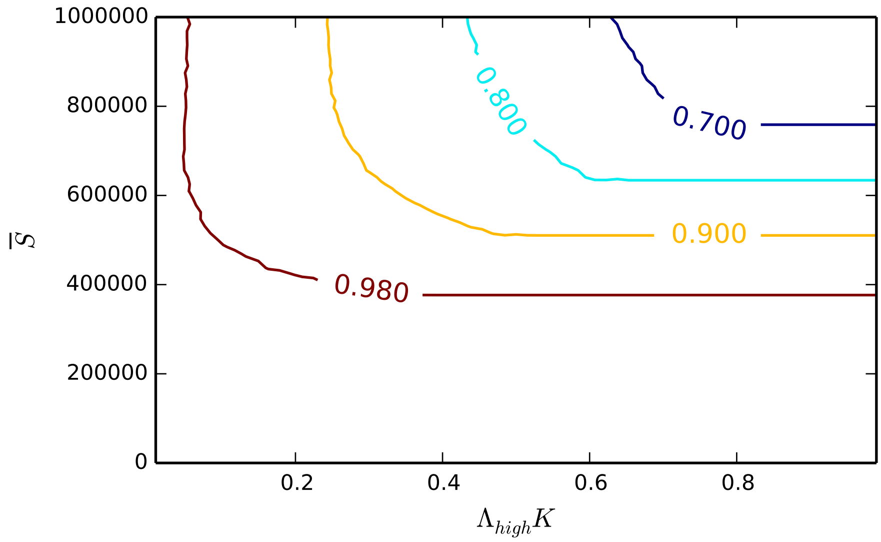

Let's now consider what happens as we vary the mix of rights, while maintaining the yield of high reliability rights (approximately) constant at 0.95.

Firstly, we see significant changes in the yield and value of low reliability rights. Further, the value of high reliability rights also varies, despite the approximately constant yield. In practice, very small changes in the allocations attached to a right (especially those received in dry years) can have a large effect on value.

The take home point from this, is that the conversion of priority rights does not work in practice. While we've adopted some simplifying assumptions here (i.e., a

Figure D.2: Constant yield combinations, low priority rights

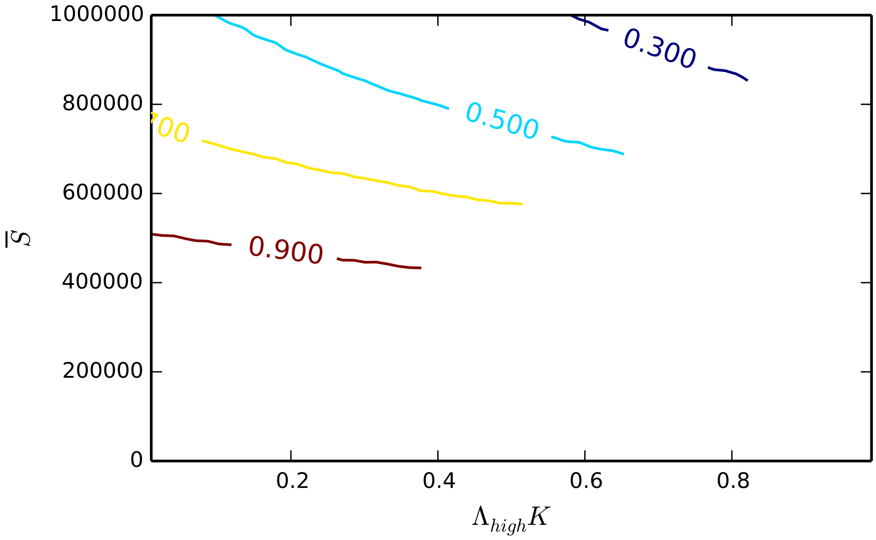

Figure D.3: Yield of low and high priority rights

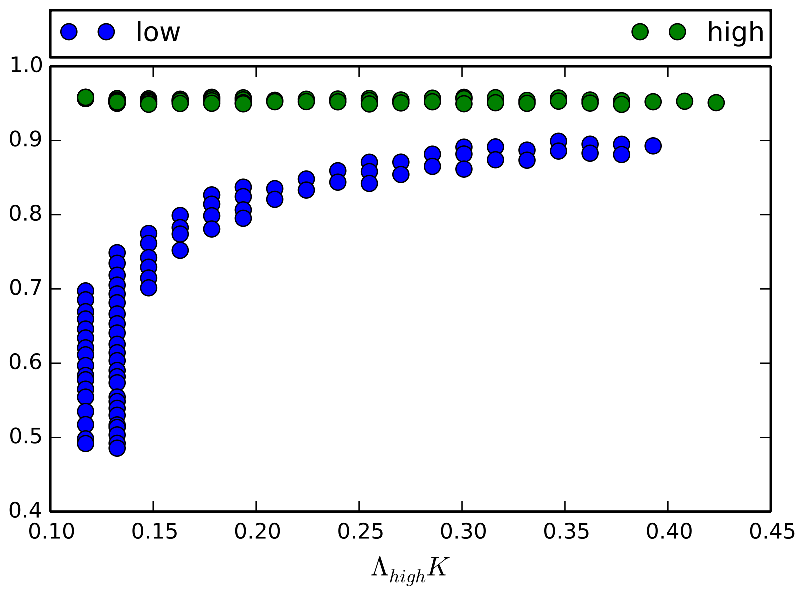

Figure D.4: Value of low and high priority rights

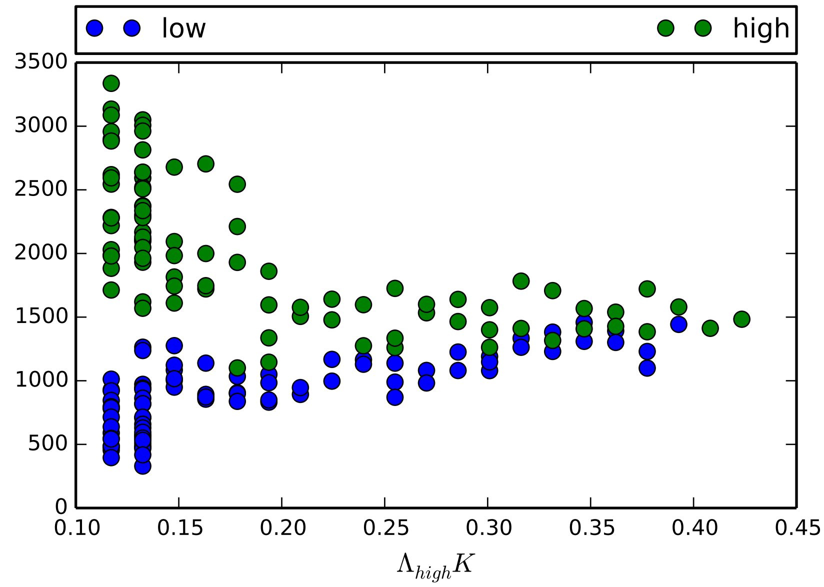

simple SOP storage rule) the idea that the value of existing rights would not be affected by conversion remains implausible.

D.3 Optimal share model

Here we solved the model 200 times for the CS-O, CS-HL-O, RS-O and RS-HL-O scenarios and recorded the final (optimal) Λ_{high} values. We then built an optimal share model by regressing the Λ_{high} values against the main parameters — using the random forest method — parameter importances are shown in table D.1.

Figure D.5 show histograms for the optimal flow shares. As we observed in chapter 6 optimal flow shares are lower in the HL scenarios. Figure D.6 plots the optimal shares against the number of high reliability users n_{high}.

As would be expected Λ_{high} is strongly correlated with the number of high reliability users n_{high} and their share of target water use $\bar{Q}_{high}/\bar{Q}$. Here $\bar{Q} = \sum_{\mathcal{U}} q_{it}$ and $\bar{Q}_{high} = \sum_{\mathcal{U}_{high}} q_{it}$ for a 'mean' water year t: in which $I_t = E[I_t]$ and $W_t = I_t$.

Λ_{high} also depends on the degree of risk aversion (higher risk aversion implies lower Λ_{high} especially under the HL scenarios).

Figure D.5: Optimal flow share Λ_{high}, histograms of 200 model runs

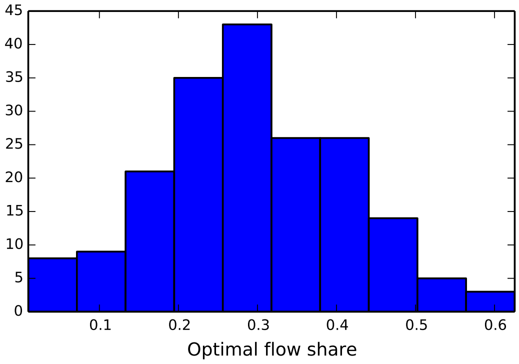

(a) RS-O

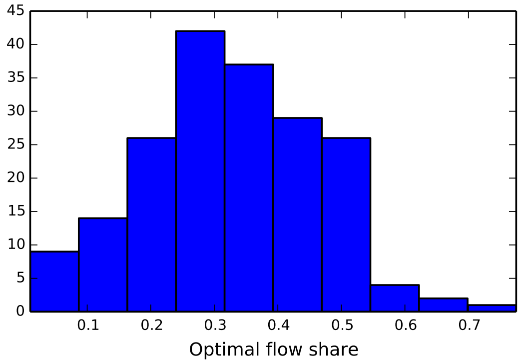

(b) CS-O

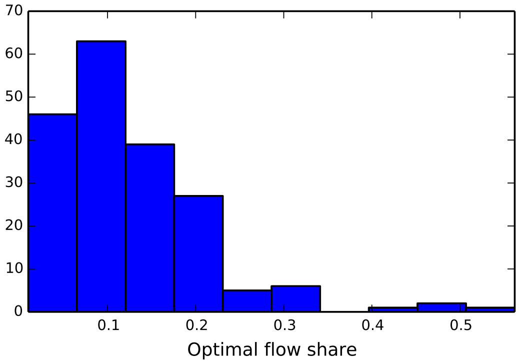

(c) RS-HL-O

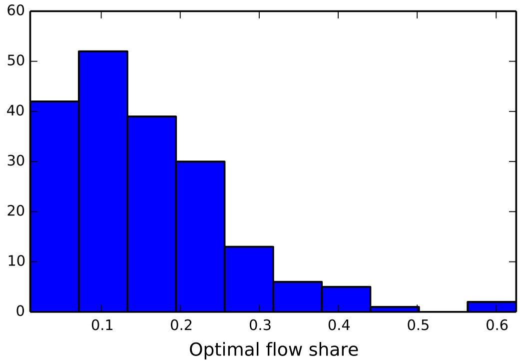

(d) CS-HL-O

Table D.1: Random forest model parameter importances

	Importance
n_{high}	37.95
$\bar{Q}_{high}/\bar{Q}$	27.31
ψ	10.27
$E[I_t]/K$	6.34
δ_{1a}	3.61
c_v	3.33
τ	3.19
ρ_I	2.92
δ_{1b}	2.43
α	1.32

Figure D.6: Optimal flow share Λ_{high} against n_{high}

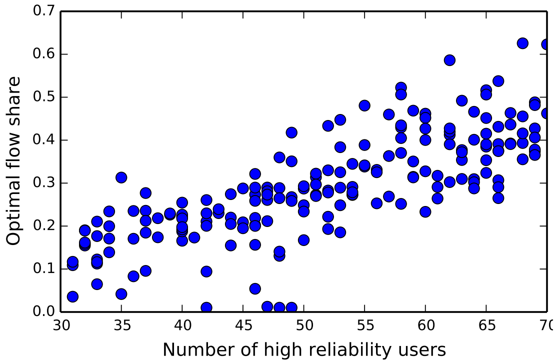

(a) RS-O

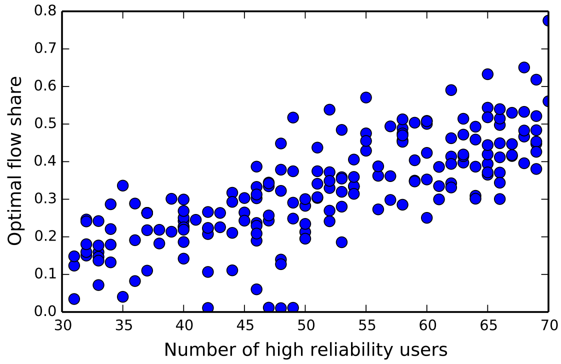

(b) CS-O

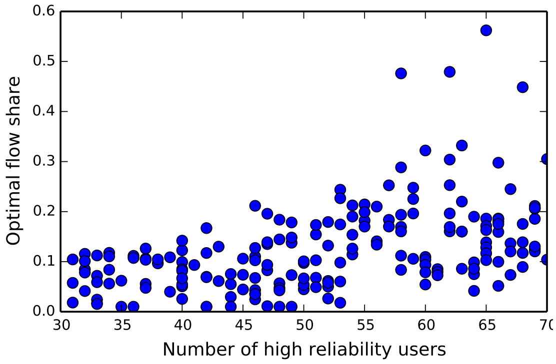

(c) RS-HL-O

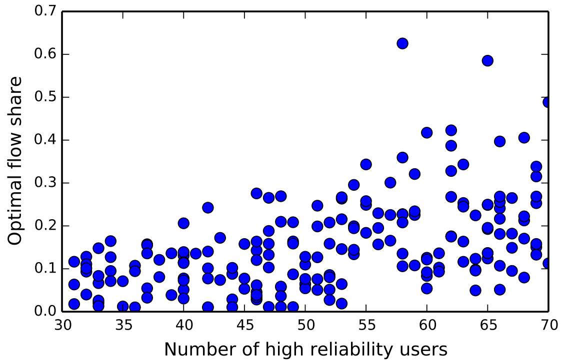

(d) CS-HL-O

Appendix E

More on environmental flows

E.1 Spot market equilibrium

The spot market equilibrium conditions differ slightly from our previous version of the model (see chapter 4.2). With $p_{it} = d_i^{-1}(a_{it}, .)$ the clearing price now satisfies the following (for all consumptive users)

$$q_{it} = \begin{cases} d_i(P_t - \tau/2, .) & \text{if } p_{it} \leq P_{it} - \tau/2 \\ a_{it} & \text{if } P_t - \tau/2 < p_{it} < P_{it} + \tau/2 \\ d_i(P_t + \tau/2, .) & \text{if } p_{it} \geq P_{it} + \tau/2 \end{cases}$$

For the EWH we have the same condition only we replace P_t with $P_t + p_\Delta$ where p_Δ is an adjustment to ensure that the EWH's budget constraint holds[1]. p_Δ is determined endogenously during the learning algorithm via a method similar to that used to find optimal inflow shares in chapter 6.

Consumptive demand

Again the consumptive users inverse demand functions d_i^{-1} are defined

$$d_i^{-1}(q_{it}, \tilde{I}_t, e_{it}) = \max\left\{\frac{\partial \pi_h(q_{it}, \tilde{I}_{it}, e_{it})}{\partial q_{it}}, 0\right\}$$

So in summer the consumptive users have the same linear demand functions as before (see section 3.5), and in winter they have zero demand.

[1] Note that we don't apply p_Δ to the EWH's payoff function, only in the spot market equilibrium conditions.

EWH demand: summer

For the EWH we define d_0^{-1} as

$$d_0^{-1}(q_{0t},.) = \max\left\{\frac{\partial B(.)}{\partial q_{0t}}, 0\right\}$$

In summer, the EWH's demand function reflects the marginal benefits of reducing extraction and increasing river flows at node three. That is, since we assume $b_2 = 0$ and we know $\frac{\partial F_{1t}}{\partial q_{0t}} = 0$ (extraction has no effect on upstream flows) we have

$$d_0^{-1}(q_{0t},.) = \max\left\{\frac{\partial B(.)}{\partial F_{3t}}\frac{\partial F_{3t}}{\partial q_{0t}}, 0\right\}$$

Now assuming $F_{2t} \geq \delta_{at}$ (given our minimum flow rule) we can define F_{3t} in terms of q_{0t}

$$F_{3t} = Z_t + q_{0t}\frac{(1-\delta_R)}{(1-\delta_{Eb})} + \bar{E}_t\delta_R$$

That is, end of system flows are just spills plus environmental flows plus return flows. Now d_0^{-1} is

$$d_0^{-1}(q_{0t},.) = \max\left\{\frac{\partial B(.)}{\partial F_{3t}}\frac{(1-\delta_R)}{(1-\delta_{Eb})}, 0\right\}$$

Here the term $\frac{(1-\delta_R)}{(1-\delta_{Eb})}$ is just the 'exchange rate' between water at the demand node and river node three.

Differentiating the benefit function we get

$$\begin{aligned}\frac{\partial B(.)}{\partial F_{3t}} &= -b_{\$}b_3e_{0t}G_I(I_t)\frac{\partial \Delta F_{3t}}{\partial F_{3t}} \\ &= \frac{2b_{\$}b_3e_{0t}G_I(I_t)}{\tilde{F}_{3t}}(1-\frac{F_{3t}}{\tilde{F}_{3t}})\end{aligned}$$

assuming that $F_{3t} < 2\tilde{F}_{3t}$ and $\tilde{F}_{3t} > 0$.

Putting this back into the inverse demand function we have

$$P_t = \hat{b}_3\hat{\delta}_0\left(1-\frac{F_{3t}}{\tilde{F}_{3t}}\right)$$

where $\hat{b}_3 = \frac{2b_\$ b_3 e_{0t} G_I(I_t)}{\tilde{F}_{3t}}$ and $\hat{\delta}_0 = \frac{(1-\delta_R)}{(1-\delta_{Eb})}$.

So we have a linear water demand function, with a maximum willingness to pay of $\hat{b}_3\hat{\delta}_0$ when $F_{3t} = 0$, declining to zero as F_{3t} approaches $\tilde{F}_{3t}$. Now inverting this we have

$$F_{3t} = \tilde{F}_{3t}(1 - P_t \hat{b}_3^{-1}\hat{\delta}_0^{-1})$$

Given our minimum flow rule F_{3t} can be defined as[2]

$$F_{3t} = Z_t + q_{0t}\frac{(1-\delta_R)}{(1-\delta_{Eb})} + \bar{E}_t\delta_R$$

we can define the demand for water q_{0t} as

$$\begin{aligned} q_{0t} &= \hat{\delta}_0^{-1}\left(\tilde{F}_{3t}(1 - P_t\hat{b}_3^{-1}\hat{\delta}_0^{-1}) - \bar{E}_t\delta_R - Z_t\right) \\ &= \hat{d}_{a0} + \hat{d}_{b0}P_t \end{aligned}$$

where $\hat{d}_{a0} = \hat{\delta}_0^{-1}\left(\tilde{F}_{3t} - \bar{E}_t\delta_R - Z_t\right)$ and $\hat{d}_{b0} = -\hat{\delta}_0^{-2}\hat{b}_3^{-1}\tilde{F}_{3t}$

EWH demand: winter

In winter there is no extraction, so all EWH purchases contribute to river flows at both nodes 1 and 3. So we have

$$\begin{aligned} d_0^{-1}(q_{0t}, .) &= \max\left\{\frac{\partial B(.)}{\partial F_{1t}}\frac{\partial F_{1t}}{\partial q_{0t}} + \frac{\partial B(.)}{\partial F_{3t}}\frac{\partial F_{3t}}{\partial q_{0t}}, 0\right\} \\ &= \max\left\{\frac{\partial B(.)}{\partial F_{1t}}\frac{1}{(1-\delta_{Eb})} + \frac{\partial B(.)}{\partial F_{3t}}\frac{1}{(1-\delta_{Eb})}, 0\right\} \\ &= \max\left\{\frac{\partial B(.)}{\partial F_{1t}}\hat{\delta}_1 + \frac{\partial B(.)}{\partial F_{3t}}\hat{\delta}_1, 0\right\} \end{aligned}$$

where $\hat{\delta}_1 = \frac{1}{(1-\delta_{Eb})}$. So we have

[2] If storage volumes are too small to supply the minimum flow, releases and allocations are all zero and the spot market is closed.

$$d_0^{-1}(q_{0t},.) = \max\left\{\hat{b}_1\left(1-\frac{F_{1t}}{\tilde{F}_{1t}}\right)\hat{\delta}_1 + \hat{b}_3\left(1-\frac{F_{3t}}{\tilde{F}_{3t}}\right)\hat{\delta}_1, 0\right\}$$

Assuming the marginal value of environmental flows is positive, we can write

$$(\hat{b}_1+\hat{b}_3) - P_t\hat{\delta}^{-1} = \hat{b}_1\left(\frac{F_{1t}}{\tilde{F}_{1t}}\right) + \hat{b}_3\left(\frac{F_{3t}}{\tilde{F}_{3t}}\right)$$

In summer F_{1t} and F_{3t} are defined as

$$F_{1t} = Z_t + q_{0t}\hat{\delta}_1 + 2\delta_{at}$$
$$F_{3t} = Z_t + q_{0t}\hat{\delta}_1$$

So we can write

$$\begin{aligned}
(\hat{b}_1+\hat{b}_3) - P_t\hat{\delta}^{-1} &= \hat{b}_1\frac{\left(Z_t + q_{0t}\hat{\delta}_1 + 2\delta_{at}\right)}{\tilde{F}_{1t}} + \hat{b}_3\frac{\left(Z_t + q_{0t}\hat{\delta}_1\right)}{\tilde{F}_{3t}} \\
(\hat{b}_1\tilde{F}_{1t}^{-1} + \hat{b}_3\tilde{F}_{3t}^{-1})(Z_t + q_{0t}\hat{\delta}_1) &= (\hat{b}_1+\hat{b}_3) - 2\delta_{at}\hat{b}_1\tilde{F}_{1t}^{-1} - P_t\hat{\delta}^{-1} \\
q_{0t} &= \hat{\delta}_1^{-1}\left(\frac{(\hat{b}_1+\hat{b}_3 - 2\delta_{at}\hat{b}_1\tilde{F}_{1t}^{-1} - P_t\hat{\delta}^{-1})}{\hat{b}_1\tilde{F}_{1t}^{-1} + \hat{b}_3\tilde{F}_{3t}^{-1}} - Z_t\right) \\
q_{0t} &= \hat{d}_{a1} + \hat{d}_{b1}P_t
\end{aligned}$$

where $\hat{d}_{a1} = \hat{\delta}_1^{-1}\left(\frac{\hat{b}_1+\hat{b}_3-2\delta_{at}\hat{b}_1\tilde{F}_{1t}^{-1}}{\hat{b}_1\tilde{F}_{1t}^{-1}+\hat{b}_3\tilde{F}_{3t}^{-1}} - Z_t\right)$ and $\hat{d}_{b1} = \frac{-1}{\hat{\delta}_1^2(\hat{b}_1\tilde{F}_{1t}^{-1}+\hat{b}_3\tilde{F}_{3t}^{-1})}$.

E.2 Central case seasonal results

Table E.1: Storage — summer opening, S_t (GL)

	Mean	SD	2.5th	25th	75th	97.5th
Planner	642.01	264.67	167.95	424.49	891.24	1,000.00
CS	590.20	252.42	160.31	389.39	801.66	1,000.00
SWA	637.49	260.50	178.70	419.82	879.38	1,000.00
OA	647.01	268.96	174.30	420.44	918.00	1,000.00
NS	514.84	259.12	134.94	301.68	706.96	1,000.00
CS-HL	570.72	255.09	140.79	371.40	776.72	1,000.00
SWA-HL	599.36	247.36	171.57	400.69	803.35	1,000.00

Table E.2: Storage — winter opening, S_t (GL)

	Mean	SD	2.5th	25th	75th	97.5th
Planner	552.92	308.96	74.40	273.96	835.42	1,000.00
CS	497.63	267.63	98.27	283.82	680.36	1,000.00
SWA	545.36	282.81	117.41	304.66	777.65	1,000.00
OA	560.53	296.59	110.73	303.05	836.92	1,000.00
NS	444.98	271.16	90.17	227.42	612.41	1,000.00
CS-HL	472.81	279.45	76.64	247.96	663.96	1,000.00
SWA-HL	503.56	268.26	106.76	287.12	688.47	1,000.00

Table E.3: Social welfare — summer, ($m)

	Mean	SD	2.5th	25th	75th	97.5th
Planner	201.50	28.06	146.56	189.74	218.00	251.60
CS	196.51	32.72	125.57	177.29	218.31	250.69
SWA	196.77	31.96	121.37	181.48	216.48	248.85
OA	195.71	31.07	119.40	181.85	214.14	244.45
NS	191.91	37.23	100.12	171.01	217.54	249.35
CS-HL	198.38	31.87	126.97	180.77	219.68	249.02
SWA-HL	197.20	32.19	129.12	179.07	218.24	251.75

Table E.4: Social welfare — winter, ($m)

	Mean	SD	2.5th	25th	75th	97.5th
Planner	15.54	18.20	0.00	0.36	27.40	60.33
CS	15.44	15.16	0.37	3.61	22.82	55.14
SWA	15.78	15.94	0.39	3.82	22.71	58.76
OA	16.26	17.34	0.34	3.19	23.64	63.72
NS	17.41	15.05	0.73	5.54	25.52	55.18
CS-HL	15.41	14.95	0.27	3.82	22.79	54.19
SWA-HL	15.05	15.35	0.35	3.17	22.32	55.79

Table E.5: Extraction — summer, 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 E.6: Extraction — winter, E_t (GL)

	Mean	SD	2.5th	25th	75th	97.5th
Planner	0.00	0.00	0.00	0.00	0.00	0.00
CS	0.00	0.00	0.00	0.00	0.00	0.00
SWA	0.00	0.00	0.00	0.00	0.00	0.00
OA	0.00	0.00	0.00	0.00	0.00	0.00
NS	0.00	0.00	0.00	0.00	0.00	0.00
CS-HL	0.00	0.00	0.00	0.00	0.00	0.00
SWA-HL	0.00	0.00	0.00	0.00	0.00	0.00

Table E.7: Withdrawal — summer, W_t (GL)

	Mean	SD	2.5th	25th	75th	97.5th
Planner	474.69	166.64	164.91	376.95	589.29	852.22
CS	494.66	211.29	140.95	324.35	648.89	886.64
SWA	474.76	162.54	135.25	357.79	611.52	695.93
OA	447.46	137.76	133.43	359.31	560.83	615.76
NS	475.89	232.32	113.55	281.84	667.76	881.48
CS-HL	497.49	201.18	140.32	336.76	674.54	824.25
SWA-HL	496.35	199.78	145.75	342.07	648.44	839.58

Table E.8: Withdrawal — winter, W_t (GL)

	Mean	SD	2.5th	25th	75th	97.5th
Planner	118.35	108.69	0.00	13.80	214.46	334.00
CS	128.96	84.32	40.26	57.49	185.75	302.28
SWA	119.28	72.39	40.26	58.31	173.32	286.53
OA	119.71	71.07	40.26	54.60	176.03	272.67
NS	157.39	80.59	40.81	87.32	216.94	300.56
CS-HL	125.06	82.63	40.26	50.52	197.82	272.56
SWA-HL	126.27	88.39	40.26	51.25	182.32	304.29

E.3 Consumptive-environmental trade-off curves

Here we solve the model with central case parameters but with varying environmental shares $\lambda_0 \in [0.1, 0.2, 0.3, 0.4, 0.5]$. The results for the CS, CS-HL, SWA, SWA-HL, OA and NS scenarios are summarised in tables E.9 to E.12. For discussion of these results refer to section 7.7.1.

Table E.9: Mean social welfare ($m)

	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

Table E.10: Mean environmental benefits ($m)

	10	20	26.3	30	40	50
CS	24.72	30.07	33.51	34.73	37.98	40.34
SWA	24.88	31.18	32.85	34.45	38.15	40.41
OA	26.72	31.74	34.17	35.36	53.03	53.07
NS	25.02	32.08	36.26	39.25	43.95	48.74
CS-HL	21.86	27.11	29.92	31.36	35.06	39.24
SWA-HL	24.58	30.07	33.06	35.22	38.83	40.53

Table E.11: Mean irrigation profits ($m)

	10	20	26.3	30	40	50
CS	187.34	182.65	178.44	176.64	170.10	161.99
SWA	184.11	182.60	179.70	177.03	170.59	163.72
OA	186.11	181.16	177.80	176.00	85.47	92.73
NS	184.54	178.39	173.05	168.69	159.63	141.17
CS-HL	189.55	186.23	183.87	182.13	178.88	173.29
SWA-HL	187.04	182.33	179.19	176.51	169.98	163.24

Table E.12: Mean net environmental water trade ($m)

	10	20	26.3	30	40	50
CS	0.02	0.01	-0.08	-0.01	-0.23	-0.15
SWA	0.01	-0.11	-0.05	-0.20	-0.15	0.07
OA	-0.01	-0.07	-0.11	-0.67	0.18	0.03
NS	0.04	0.03	0.11	-0.00	0.02	0.04
CS-HL	0.01	0.01	0.05	0.01	0.06	0.14
SWA-HL	-0.01	-0.04	-0.11	-0.01	-0.16	-0.08

Table E.13: Mean storage, S_t (GL)

	10	20	26.3	30	40	50
CS	581.44	561.37	543.91	550.13	545.51	529.00
SWA	599.06	598.54	591.43	579.59	575.40	597.67
OA	618.76	614.59	603.77	603.55	894.61	923.44
NS	509.00	494.05	479.91	474.49	452.62	405.78
CS-HL	538.53	527.05	521.77	521.76	524.29	519.46
SWA-HL	556.09	553.57	551.46	546.68	561.08	553.81

E.4 Central case user group results

User group results

Table E.14: Low reliability user payoff, $\sum_{\mathcal{U}_{low}} u_{it}$ (GL)

	Mean	SD	2.5th	25th	75th	97.5th
Planner	0.00	0.00	0.00	0.00	0.00	0.00
CS	74.52	12.93	52.76	63.45	86.80	94.67
SWA	74.64	12.27	52.54	64.70	85.81	93.22
OA	73.08	10.87	51.22	64.65	82.21	88.17
NS	74.11	12.24	54.82	63.74	84.96	95.06
CS-HL	77.16	19.06	40.64	62.41	93.76	103.90
SWA-HL	74.70	12.75	52.94	63.96	86.64	94.24

Table E.15: High reliability user payoff, $\sum_{\mathcal{U}_{high}} u_{it}$ (GL)

	Mean	SD	2.5th	25th	75th	97.5th
Planner	0.00	0.00	0.00	0.00	0.00	0.00
CS	103.92	17.04	49.36	100.46	112.93	119.33
SWA	105.06	17.95	44.72	101.68	114.27	121.27
OA	104.72	18.82	39.43	101.68	114.15	120.94
NS	98.94	22.38	23.91	98.32	111.07	117.76
CS-HL	106.71	12.26	82.72	105.77	111.62	115.79
SWA-HL	104.49	16.59	52.96	101.07	113.21	119.97

Table E.16: Low reliability use, $\sum_{\mathcal{U}_{low}} q_{it}$ (GL)

	Mean	SD	2.5th	25th	75th	97.5th
Planner	216.73	122.69	0.00	129.04	305.49	443.28
CS	229.11	113.33	34.27	138.11	323.36	444.71
SWA	213.63	97.68	33.92	136.20	298.44	373.01
OA	195.89	79.04	27.17	142.00	256.14	320.71
NS	214.45	110.24	34.95	127.13	299.37	436.34
CS-HL	294.22	185.74	0.00	132.91	460.06	614.22
SWA-HL	224.07	109.96	33.01	136.91	314.62	423.14

Table E.17: High reliability use, $\sum_{\mathcal{U}_{high}} q_{it}$ (GL)

	Mean	SD	2.5th	25th	75th	97.5th
Planner	131.66	26.58	90.66	113.54	152.80	182.76
CS	155.21	67.79	31.29	98.73	212.07	263.75
SWA	166.45	69.19	32.29	113.95	217.73	281.50
OA	160.98	65.53	28.31	116.97	210.44	277.58
NS	136.24	66.97	22.79	80.53	189.05	257.66
CS-HL	122.44	19.85	77.93	115.37	127.19	160.99
SWA-HL	155.38	66.51	34.91	99.34	211.30	257.15

Table E.18: Low reliability user storage, $\sum_{\mathcal{U}_{low}} s_{it}$ (GL)

	Mean	SD	2.5th	25th	75th	97.5th
Planner	0.00	0.00	0.00	0.00	0.00	0.00
CS	194.09	107.45	24.59	107.33	277.10	398.90
SWA	195.88	105.20	23.21	107.97	282.75	383.63
OA	203.20	110.65	21.30	108.38	306.25	384.06
NS	178.05	104.98	25.12	93.03	249.77	398.90
CS-HL	234.79	167.84	0.00	89.69	362.42	555.23
SWA-HL	195.76	107.89	24.89	107.71	279.98	398.90

Table E.19: High reliability user storage, $\sum_{\mathcal{U}_{high}} s_{it}$ (GL)

	Mean	SD	2.5th	25th	75th	97.5th
Planner	0.00	0.00	0.00	0.00	0.00	0.00
CS	174.94	80.90	28.05	110.17	246.84	293.33
SWA	192.41	99.10	26.31	108.03	284.57	353.61
OA	209.65	116.45	22.22	108.40	311.22	410.46
NS	130.93	77.19	18.47	68.41	183.67	293.33
CS-HL	123.46	19.21	59.37	127.37	130.06	137.00
SWA-HL	174.47	80.47	28.87	107.79	246.47	293.33

E.5 No trade scenario

Here we solve the model of chapter 7 with the central case parameters but with no spot market (with $\tau = \infty$). The results are shown in tables C.2 to C.7.

Comparing these results with the central case, we estimate the mean *gains from trade* (figure E.1). As we found in appendix E the gains from spot market trade are much higher in the absence of storage rights (in the NS scenario). However, with environmental demands the gains from trade are significant, even with well defined storage rights (in the order of \$6m a year).

In terms of storage, we see that the removal of trade generally increases mean storage volumes. OA for example now leads to substantial over-storage with a mean volume of 720 GL compared with around 600 under the planner case.

Figure E.1: Gains from trade (\$m)

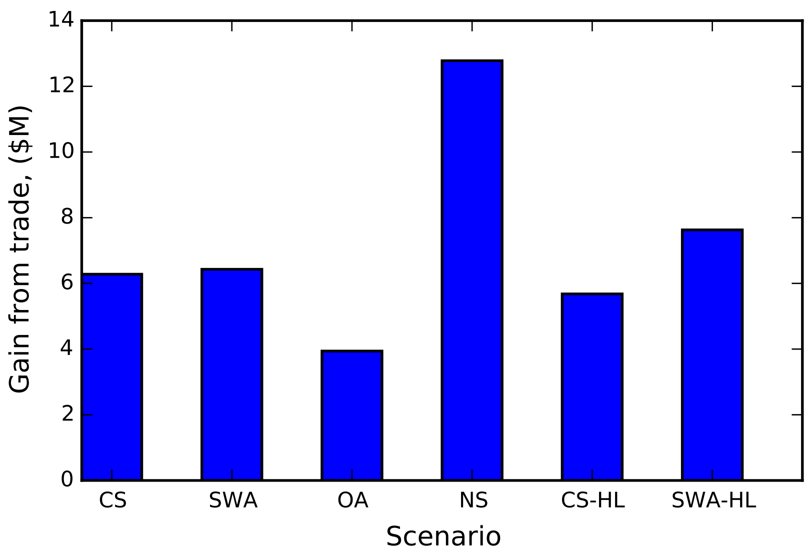

Table E.20: Social welfare, $\sum_{i=1}^{n} u_{it}$ (\$m)

	Mean	SD	2.5th	25th	75th	97.5th
Planner	216.55	38.40	149.93	192.22	240.99	294.61
CS	205.67	48.47	81.63	181.00	237.44	287.64
SWA	206.12	44.52	93.51	183.39	232.97	289.16
OA	208.03	41.58	114.05	185.83	230.57	294.15
NS	196.54	55.67	66.59	165.03	235.71	283.44
CS-HL	208.11	41.28	123.70	181.01	235.69	284.18
SWA-HL	204.62	47.88	81.01	180.87	235.71	284.04

Figure E.2: Mean profits versus environmental benefits

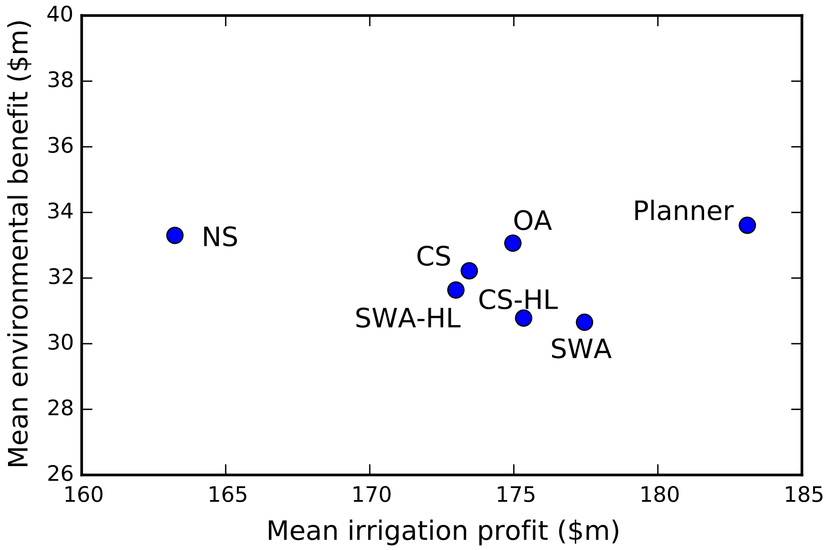

Figure E.3: Mean storage, S_t by iteration

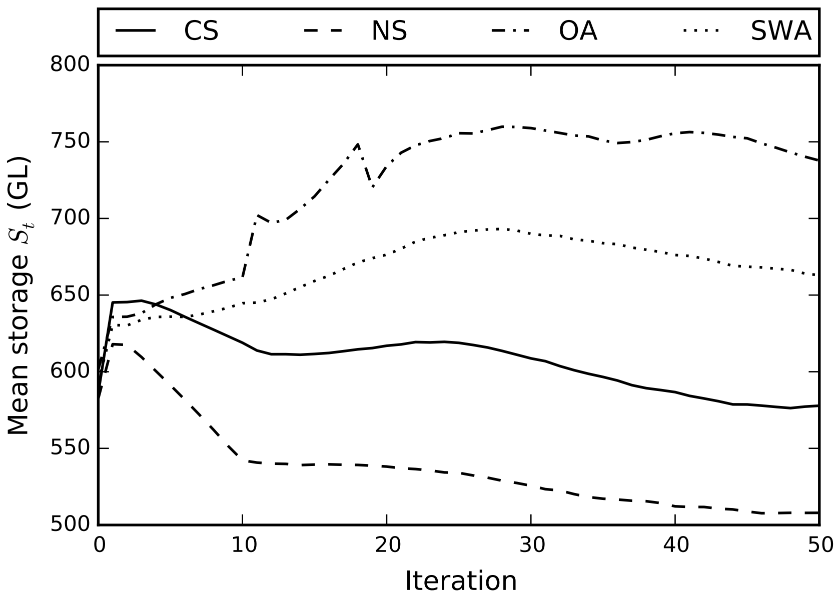

Table E.21: Storage, S_t (GL)

	Mean	SD	2.5th	25th	75th	97.5th
Planner	616.95	270.49	139.75	383.84	863.44	1,000.00
CS	576.77	250.21	142.82	375.74	781.64	1,000.00
SWA	663.89	250.42	176.41	470.99	888.62	1,000.00
OA	720.35	247.41	198.64	531.07	948.84	1,000.00
NS	503.93	254.40	113.66	298.48	696.66	1,000.00
CS-HL	519.23	243.79	134.65	326.08	693.53	1,000.00
SWA-HL	563.91	246.31	139.64	366.71	754.32	1,000.00

Table E.22: Consumptive profits ($m)

	Mean	SD	2.5th	25th	75th	97.5th
Planner	182.43	18.32	145.67	174.72	194.54	206.86
CS	173.45	31.59	73.75	166.30	192.92	204.05
SWA	175.34	27.53	85.40	169.92	191.30	201.26
OA	174.97	22.35	104.97	171.06	187.02	196.22
NS	163.24	39.00	59.83	145.69	190.65	202.09
CS-HL	177.45	23.67	117.71	165.67	193.57	206.77
SWA-HL	172.99	32.06	72.23	165.65	192.66	204.03

Table E.23: Environmental benefits ($m)

	Mean	SD	2.5th	25th	75th	97.5th
Planner	34.12	30.27	1.46	8.76	52.43	107.52
CS	32.22	25.62	1.16	12.30	46.28	96.05
SWA	30.78	26.11	1.13	11.18	43.39	99.05
OA	33.06	29.22	1.17	11.65	45.75	111.43
NS	33.30	24.44	1.81	14.16	47.44	92.43
CS-HL	30.65	24.45	2.15	11.06	43.80	91.35
SWA-HL	31.64	24.57	1.58	12.61	44.91	92.61

Glossary

capacity sharing A system of property rights for water in regulated rivers proposed by Dudley and Musgrave (1988) in which users hold shares in both storage inflows and capacity. , 9, 80, 85

carry-over rights A form of storage right adopted in the southern MDB, permitting the transfer of unused water allocations between annual accounting periods. , 257

consumptive Consumptive water use occurs when water is extracted from a river, with little if any subsequent return flow. Typically consumptive use supports private benefits. Examples include irrigation and household water use. , 13

continuous accounting A system of property rights for water adopted in some northern MDB rivers. The defining feature being inter-annual (i.e., monthly) updating of water accounts for user withdrawals, as opposed to the annual accounting of the southern MDB. , 209, 257

environmental water holder A government or other not-for-profit organisation which holds water property rights and uses them to achieve in-stream environmental benefits. A key Australian example is the Commonwealth Environmental Water Holder.

fitted Q iteration A reinforcement learning algorithm involving the estimation of a 'state-action' value or Q function. Unlike standard Q learning, the Q function is fit to a large batch of simulated samples rather than being updated after each time period. , 183

institution costs Transaction costs that are not marginal to particular transactions (fixed transaction costs). Institution costs include the costs of establishing and enforcing property rights.. , 63

internal spill Where a water user's account reaches its maximum limit, such that any further inflow credits are forfeited and reallocated to other users. , 81

irrigation district A particular type of collective organisation representing the interests of a group of irrigators within a defined geographical boundary. Common in the western US. , 73

large dam A dam greater than 15 metres in height or 3 gigalitres in volume. , 2, 22

machine learning A sub-field of computer science concerned with algorithms that 'learn' from data with minimal human input. Also known as artificial intelligence. Machine learning is typically divided into supervised learning, unsupervised learning and reinforcement learning. , 181, 187, 208, 246

Markov decision process A discrete time stochastic control process, representing the problem of an agent taking actions in a stochastic environment in order to maximise a reward. , 179

multi-agent learning An interdisciplinary field concerned with the application of reinforcement learning methods to multi-agent problems, specifically stochastic games. , 198

non-consumptive Non-consumptive or in-stream water use is where the flow of water within the river generates some — often public — benefit such as environmental improvements or flood mitigation. , 13

prior appropriation A system of property rights for water in unregulated rivers, first developed in the western US during the 1800s. A defining feature being the concept of 'first in time, first in right' where water is allocated in priority to users with the oldest water rights. , 42, 71

priority rights Water property rights where flows are allocated to certain user groups (i.e., the high reliability or 'senior' rights holders) before others (i.e., the low reliability or 'junior' rights holders). , 111

random forest A non-parametric regression method popular in supervised learning, consisting of an ensemble of decision trees. , 255

regulated river A river system where the flow is controlled to a significant extent by one or more large dams. , 6, 13

reinforcement learning A sub-field of machine learning concerned with the design of algorithms for solving Markov decision problems. Commonly reinforcement learning algorithms do not require knowledge of the 'environment': the transition or payoff functions. , 181

release sharing A term used by Dudley and Musgrave (1988) to describe property rights for water where storage decisions are made by a central agencies and users receive shares in releases. , 59

return flow Where some of the water extracted from a river for consumptive use flows back into the river system. , 14, 35

spill Where a reservoir reaches its capacity (water rises to the height of the dam wall) such that further inflows spill uncontrolled down stream. , 24

spill forfeit rules An approach to defining property rights to water storage capacity, where users face account deductions in the event of actual storage spills. , 81

stochastic game A dynamic game with stochastic state transitions or informally a multi-agent Markov decision process. , 180

storage rights Water property rights which allow users to hold private storage reservers in public reservoirs. , 77

supervised learning A sub-field of machine learning concerned with inferring a continuous function from a set of training data (i.e., non-parametric regression). , 187

tile coding A function approximation technique popular in reinforcement learning, credited to (Albus 1975). Also known as cerebellar model articulation controller or CMAC. , 190