Hydrologic-Economic Modeling of Irrigated Agriculture in the Lower Murrumbidgee Catchment: Investigations Into Sustainability

by

Christopher M. Stubbs

A.B., Physics, University of California, Berkeley (1988)
S.M., Technology and Policy, Massachusetts Institute of Technology (1996)
S.M., Environmental Engineering, Massachusetts Institute of Technology (1996)

Submitted to the Department of Civil and Environmental Engineering
in partial fulfillment of the requirements for the degree of

Doctor of Philosophy
in Hydrology and Water Resources Engineering

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

June 2000

Author .....
Department of Civil and Environmental Engineering
May 10, 2000

Certified by .....
Dennis B. McLaughlin
H.M. King Bhumibol Professor of Water Resource Management
Thesis Supervisor

Accepted by .....
Daniele Veneziano
Chairman, Department Committee on Graduate Students

Hydrologic-Economic Modeling of Irrigated Agriculture in the Lower Murrumbidgee Catchment: Investigations Into Sustainability

by

Christopher M. Stubbs

Submitted to the Department of Civil and Environmental Engineering
on May 10, 2000 in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy in Hydrology and Water Resources Engineering

Abstract

Increasing water scarcity and growing demand for food have made better management of land and water resources essential to maintaining the sustainability of irrigated agriculture. Policies designed to improve environmental quality and irrigated production need to be analyzed in an integrated framework. We present a catchment-scale hydrologic-economic model of irrigated agriculture which is dynamic and spatially distributed. It can be used to evaluate land and water policies designed to manage irrigation-induced salinization.

The model incorporates hydrologically realistic representations of groundwater flow and soil salinization into an economic optimization framework. The sum of discounted net revenues from irrigation over the planning horizon is maximized by choosing annual areas planted to each crop in each of the economic subregions. The groundwater system is represented using a linear state-space model derived from a finite-difference approximation of the groundwater flow equation. The number of groundwater states is substantially reduced using balanced truncation, a technique used in control engineering. A simple representation of the salinization process is derived from detailed numerical simulations of unsaturated zone flow and salt transport. These detailed simulations include realistic meterological forcing, crop root extraction, and the effect of shallow, saline watertables.

The use of the model for policy analysis is demonstrated in a case study of the Lower Murrumbidgee Catchment. The study area is in the Murray-Darling Basin of Australia and includes a major irrigation district threatened by salinization from rising watertables. We first simulate socially optimal management over a 15-year planning horizon. The socially optimal solution internalizes the externalities of the common-pool groundwater system and allows redistribution of water allocations to different areas. This solution is compared to scenarios which include the common-pool externality and policy options in various combinations. The policy options considered are a restriction on the amount of cropland planted to rice and the trading of surface water allocations. We find the rice area restriction decreases economic net benefits while water trading increases net benefits. There is little difference between the social optimum and the common-pool scenarios suggesting that the cost of the common-pool externality is small.

Thesis Supervisor: Dennis B. McLaughlin
Title: H.M. King Bhumibol Professor of Water Resource Management

Acknowledgments

First of all, I would like to thank my advisor, Dennis McLaughlin, for giving me the opportunity to tackle this project and for helping to guide me through its completion. Useful suggestions were also provided by the other members of my doctoral committee: Charles Harvey and Fatih Eltahir from MIT, and Richard Howitt from UC Davis. This work was supported by the MIT/ETH/UT Alliance for Global Sustainability.

Many people contributed to the work on the Australian case study. GIS data was provided by the Australian Surveying and Land Information Group and by Tim Gates of Colorado State University. Scott Lawson and Ary van der Lely, both of NSW DLWC in Leeton, were very generous with their time when answering my questions and providing data. The groundwater model of the Lower Murrumbidgee was provided by Charles Demetriou of the NSW DLWC in Paramatta.

The staff at CSIRO Land and Water in Griffith has been very supportive of our research efforts. In particular, John Blackwell was a gracious host during our initial visit to Australia. Special thanks to Shabaz Khan for his interest in our work and for providing helpful feedback during his visit to MIT.

I would like to thank Bruce Sutton for inviting me to the University of Sydney as a visiting scholar. Sailing the Sydney harbor and exploring the beaches gave me a whole new perspective on things. Many thanks to Tim and Lynn Clune for their hospitality and willingness to teach us Australian English.

I am grateful to Prof. Wolfgang Kinzelbach at the ETH-Zurich for inviting me to his institute and hosting some spirited discussions. Thanks also to Kai Zoellmann and Joerg Schulla for their collaboration.

A big thank you to the McLaughlin Group, past and present, and everyone else in the Parsons Lab. Freddi has been a great friend from the very beginning. Rolf has been both a loyal lunch companion and friend. Kirsten has helped keep my karma on track. Julie has been a pleasure to swap GIS tips with. Lynn, Feng and Cheng-Ching were always amusing computer room companions.

This thesis is dedicated to my lovely wife Sandy. Thanks for your love and support, and for being such a wonderful partner on this adventure.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

Land under irrigation produces over one-third of the global food supply while occupying only around 17% of the world's cropland [Hillel, 1991]. Irrigated agriculture will clearly serve a very significant role in feeding the growing human population. Its sustainability is an important issue which has been considered in several recent works for both general audiences [e.g., Postel, 1999; Hillel, 1991] and for academics [e.g., Letey, 1994; van Schilfgaarde, 1996; Rhoades, 1997]. These works have brought attention to the many threats to the sustainability of irrigation. However, they all conclude that the technical means exist to sustain irrigation as long as we better manage our land and water resources. One of the most serious issues which requires management is soil salinization. In this thesis we develop a model of soil salinization at a regional scale which can be used to analyze both the economic and hydrologic effects of management policies.

1.1 Irrigation-Induced Soil Salinization

The extent of irrigation-induced salinization has been examined in several studies, but is known only very roughly. The World Bank [1992] estimated that salinization caused by irrigation affects about 60 million ha worldwide. This corresponds to 24% of all irrigated land. The problem mainly occurs in arid and semi-arid regions. On the country scale, it is estimated that in Australia 20 percent of irrigated land is affected by salinity; in Pakistan, 14 percent; in Israel, 13 percent; in China, 15 percent; in Egypt, 30 percent; and in Iraq, 50 percent [Gleick, 1993; Ghassemi et al., 1996].

1.1.1 Physical Process of Salinization

Salinization is an increase of soluble salt concentrations in soil water, surface waters or groundwaters. We focus our attention on soil salinization, which is particularly a concern in agricultural areas where crop yields can be reduced by soil salinity levels above certain thresholds. Plants have different levels of tolerance to soil salinity, but at high enough levels all plants will be affected. Saline soils occur naturally in many regions, but their effect on

plants is often exacerbated by irrigation, which dramatically changes the water and salt balance in the soil and groundwater.

The salt balance of the rootzone is what determines the extent of salinity damage to crops [Marshall et al., 1996]. All irrigation waters contain dissolved salts. When this irrigation water is applied to the soil, the water is removed by evapotranspiration (ET) leaving the salt behind in the rootzone. Salt may also originate from soil weathering or be transported by rain or dust. In order for irrigation to continue, this salt must be leached down from the rootzone by water in excess of the crop requirements. If it is not, salt will build up in the rootzone. See Figure 1.1. This is not an issue in humid regions because natural rainfall is sufficient to provide leaching. In more arid areas, however, this leaching must be done using irrigation water.

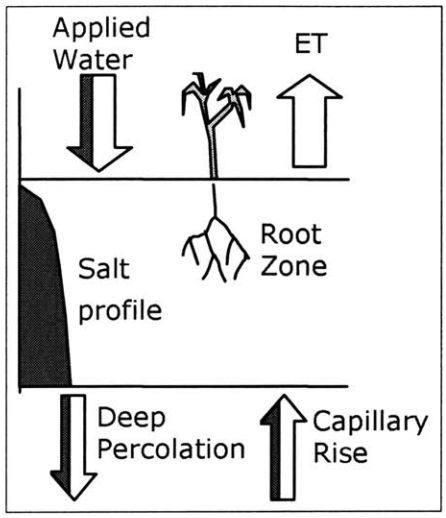

Figure 1.1: Conceptual Rootzone Salt Balance
Applied water contains salt, but ET does not. To prevent salt accumulation in the rootzone, the downward salt flux in the deep percolation must exceed the upward flux from capillary rise.

Salt that is leached from the rootzone is stored in the lower soil profile or in the groundwater. If there is more than several meters of unsaturated zone separating the rootzone and the stored salts, the salt can never travel upwards back into the rootzone as long as the groundwater levels do not rise. In many cases, however, deep percolation from irrigation results in a rise in groundwater levels. When groundwater levels are close to the ground surface, there are two effects which have a tendency to increase rootzone salinity levels: 1) salt stored in the lower soil profile and groundwater can now be pulled into the rootzone by evapotranspiration, and 2) the hydraulic head gradient driving water downward during leaching events is reduced so that leaching is less effective. As a general guideline, once the watertable rises to within 2 meters of the ground surface, there is a danger of salinization.

1.1.2 Engineering Solutions to Salinization

The groundwater system has a certain natural capacity to dissipate deep percolation water. While some amount of leaching is required to maintain crop yields, leaching beyond this may be harmful if it exceeds the groundwater system capacity and causes watertables to rise too high. In semi-arid regions, this capacity is often quite small under natural conditions. It can be increased, however, through the installation of expensive surface or subsurface drainage systems. Drainage systems remove excess rain or irrigation water before it has a chance to enter the groundwater system. In addition, the drainage system can remove enough groundwater to prevent watertables from rising high enough to endanger crops. The installation of artificial drainage may not be economically feasible for protecting low-value crops or pastures. In addition, there may be environmental problems disposing of the drainage water [National Research Council, 1989].

Excess leaching can also be reduced through better irrigation management and the use of modern pressurized irrigation systems. Other engineering approaches for drainage control or disposal include lining canals to reduce leakage, building evaporation basins for drainage disposal, and reusing drainage water on salt tolerant crops.

1.1.3 Economic Analysis of Salinization

There are often also economic causes of salinization. These arise because two essential resources, irrigation water and the assimilative capacity of the groundwater system, are not priced or allocated correctly to reflect their scarcity values and opportunity costs [Wichelns, 1999]. Irrigation water is typically priced well below its economic value and farmers can discharge drainage water to the groundwater system with no charge or restriction. These factors provide incentives for farmers to over-irrigate and under-invest in efficient irrigation and drainage systems.

These economic causes of salinization can be addressed through policy changes which correct the economic incentives. For example, by raising the price of water to its true economic price or introducing water trading schemes, farmers have an incentive to use water more efficiently. Other policies may rely on regulations which restrict groundwater recharge or crop choice. In many cases, implementing these policy changes may be less expensive than large engineering works.

1.2 Previous Work

There is an extensive literature concerned with the environmental impacts of agriculture. Only a few of these studies combine a physically realistic representation of environmental processes with an economically realistic description of the production process at regional scales. The resulting models (which we call "hydrologic-economic" models) are typically formulated as optimization problems [Taylor and Howitt, 1993]. In these models farmers are assumed to be profit maximizing as a first approximation, so profit or net revenue is the objective function of the optimization.

One of the first large-scale hydrologic-economic models was a model of the San Joaquin Valley in California, described in California Department of Water Resources [1982]. The focus of this model was to assess the impact on agriculture of groundwater depletion; it did not consider salinity. A more recent model of the San Joaquin which includes salinity is the Westside Agricultural Drainage Economics (WADE) model [Hachett et al., 1991]. Unlike the earlier model, it maximizes short-run rather than long-term revenue.

Another early large-scale model was a linear programming model of the Indus Basin in Pakistan developed by the World Bank and described in Bisschop et al. [1982]. This model was used to evaluate proposed water projects and agricultural policies in what is now the world's largest contiguous surface distribution system. The Indus Basin model considers agricultural production and consumption, irrigation infrastructure, and groundwater quality and depth. The model was recently revised and used to analyze waterlogging and salinity problems in the Indus Valley [Ahmad and Kutcher, 1992].

Lefkoff and Gorelick [1990] developed a "Hydrologic-Economic-Agronomic" model which accounts explicitly for water quality degradation related to salinization. The model was used to evaluate the benefits of a water rental market in a stream-aquifer system in southeast Colorado. By maximizing the short-run profits of each farm individually, it represents the externality effects of saline drainage. It considers groundwater flow, transport of dissolved solids, stream-aquifer interactions, and irrigation with saline water. It does not consider soil salinity. Groundwater flow and salt transport are described with simplified constraints derived using a response matrix approach.

Lee and Howitt [1996] used a basin-scale model to evaluate policies which address salinity externalities from irrigated agriculture in the Colorado River. This model considers crop activities in multiple regions and it includes a detailed set of hydrologic constraints and an objective function which accounts for agricultural returns, salinity control costs, and water quality benefits. It does not consider groundwater. The hydrologic portions of the model are described in Lee et al. [1993].

Shah et al. [1995] examine the difference between optimal and common property solutions to agricultural drainage problem using an exhaustable resource framework. Their approach is similar to ours in that they use a dynamic and integrated framework, but it is not spatially distributed.

The agricultural production aspects of a hydrologic-economic model are typically described by production functions which relate crop yield to inputs and other factors, such as the quantity and quality of applied irrigation water, irrigation technology, and environmental conditions. Crop production functions can be classified as seasonal or transient [Letey et al., 1990]. Seasonal production functions relate seasonal applied water quantity and quality to seasonal yield. Letey et al. [1990] and Letey and Dinar [1986] describe seasonal models which include salinity. Transient production functions are derived from detailed simulation models of the unsaturated zone and crop growth based on Richard's equation with a root extraction term.

1.3 Need For an Integrated Framework

In this research, we focus our analysis on land and water policies as a means of managing regional-scale salinization. This analysis requires an integrated framework including both the biophysical and economic systems. A change in the price of irrigation water or the imposition of cropland use restrictions will change the crop mix and water use over the region. How the farmers in each part of the region respond will depend on the specific characteristics of that part. These characteristics would include the depth to groundwater and the predominant soil types. Farmers can respond to rising watertables by altering their crop mix and investing in more efficient irrigation technologies.

Through the groundwater system, the cropping decisions of farmers in different parts of the region are linked. We need to model the economic response of farmers to know how regulations will influence economic decisions such as crop choice. At the same time, we need to model the hydrologic response since this influences the economic response through groundwater levels. These linkages are the motivation for developing an integrated hydrologic-economic modeling framework.

To illustrate the framework, we apply it to the salinization issues of a major irrigated area in semi-arid Australia: the Lower Murrumbidgee Catchment. Before European settlement, the semi-arid climate and natural vegetation kept areal recharge to an insignificant level. The clearing of native deep-rooted vegetation and its replacement with intensively irrigated crop species led to a large increase in recharge to groundwater. A mound was created under the irrigation areas with watertable depths in many areas of less than 2 meters. These high watertables have caused crop yield reductions.

Past efforts to control high watertables have emphasized landuse restrictions. In particular, rice areas have been targeted in order to reduce recharge. Another policy issue is the effect of a recently introduced system of trading of surface water allocations. We will investigate these policy issues using our integrated framework.

1.4 Research Issues

The main research issues addressed in this thesis are:

- We develop an integrated, hydrologic-economic model of irrigated agriculture which is dynamic and spatially distributed. The model can be used to investigate the regional-scale effects of alternative land and water management policies on crop production, groundwater levels, and economic returns to irrigation.

- We develop hydrologically accurate and computationally feasible representations of deep percolation and crop yields under different watertable depths and salinities. We use a systems analysis technique to reduce the computation cost of representing groundwater flow.

- We use the model to evaluate policy options for managing soil salinity in a case study of the Lower Murrumbidgee Catchment, a major irrigation area in Australia. We

analyze the economic and hydrologic effects of land-use restrictions and water entitlement trading under optimal management and a common-property arrangement for the groundwater system.

Chapter 2

Lower Murrumbidgee Catchment Study Area

In this chapter we give a general description of the case study area including a summary of agriculture, water resources, and indicators of salinization. We then describe the existing institutions concerned with land and water management. The chapter concludes with a description of the policy issues which we will examine using our hydrologic-economic model.

2.1 General Description

The Murrumbidgee Catchment is part of Murray-Darling River Basin, one of the most productive agricultural areas in Australia. The Murray-Darling Basin covers 1 million square kilometers (14% of Australia), and most of Australia's irrigated agriculture is produced there. The Murrumbidgee Catchment covers around 84,000 km^2. Annual irrigated agricultural output in the catchment averages about AUS$400 million, which is around one-third of the total agricultural production. The major irrigated enterprises are rice, wheat, citrus, wine grapes, peaches, vegetables, prime lambs, wool and beef cattle.

The case study focuses on the area surrounding the Murrumbidgee Irrigation Areas and Districts (MIA), one of the oldest and most productive of Australia's irrigated areas. Figure 2.1 shows the Lower Murrumbidgee Catchment, including the location of our study area and the irrigation district boundaries. The MIA includes the Yanco and Mirrool Irrigation Areas (centered around the towns of Leeton and Griffith) and the Benerembah, Tabbita and Wah Wah Irrigation Districts. Large-scale irrigation first began in the Yanco Irrigation Area following the construction of Burrinjuck Dam in 1912.

In addition to the MIA, our study area includes nearby irrigated farms outside of the designated irrigation districts. These farms rely on water from the Murrumbidgee diverted through privately constructed irrigation canals and groundwater. Our analysis does not include the Coleambally Irrigation Area and the Lowbidgee Flood Control and Irrigation District due to data and resource limitations. Future work may extend the analysis to the entire Lower Murrumbidgee Catchment.

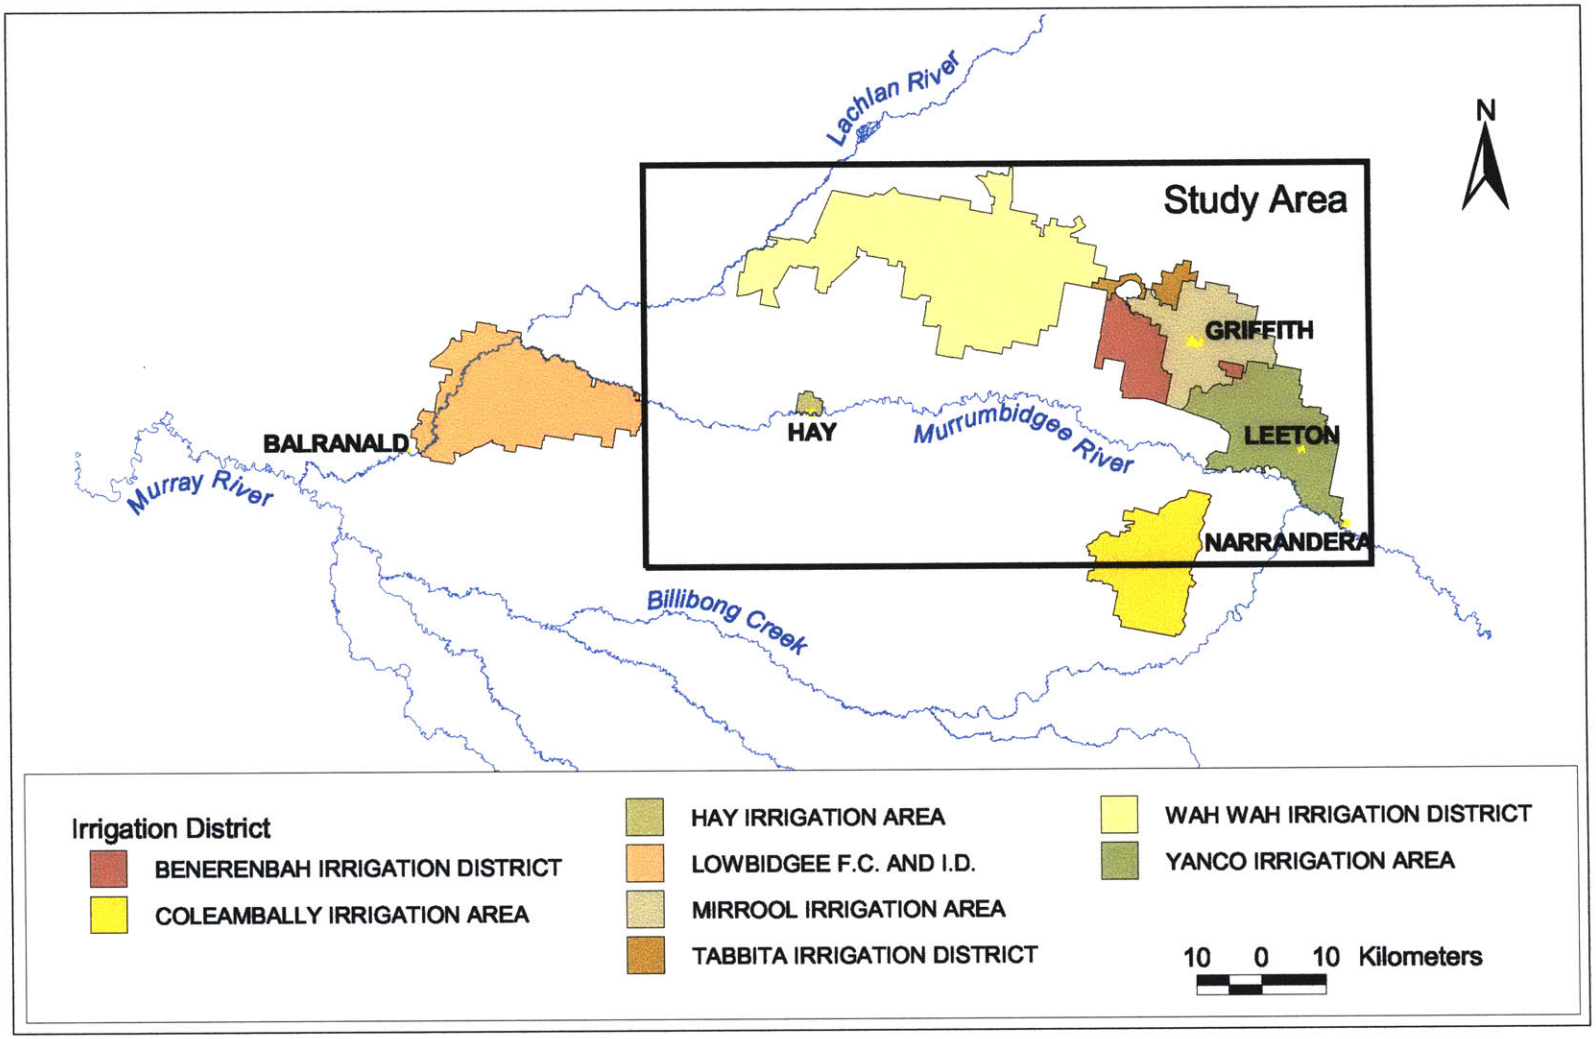

Figure 2.1: Lower Murrumbidgee Valley and Study Area

The largest town in the study area is Griffith with a population of 21,000. Other smaller towns include Leeton, Narrandera and Hay. The total population in the study area is around 36,000 [MIA LWMP Taskforce, 1998]. The regional economy is based around agriculture.

The study area is in geological transition zone. To the east, the landscape becomes increasingly hilly and Paleozoic outcrops become more common. To the west, the topography is basically flat open plains. Rock outcrops rise over 300 meters above the plains but only occur on the north-eastern edge of the study area. The elevation varies from 135 m ASL in the east to 70 m ASL in the west. Surface elevations are shown on Figure 2.2.

The climate is semi-arid with periods of flooding. The average monthly rainfall and pan evaporation at Griffith is shown in Figure 2.3. Monthly rainfall is approximately uniform over the year with an annual average of 409 mm [Australian Bureau of Meterology, 1988]. Average monthly pan evaporation varies from a low of 44 mm/month in June to a high of 297 mm/month in December. The average annual evaporation is 1,827 mm/year. Average monthly temperatures range from 8.5 C in July to 24.1 C in February. The mean daily high temperatures during the summer (December to February) are around 30 to 33 C, although temperatures in the summer may exceed 38 C. In the winter, the high temperature averages from 14 to 17 C, and the low temperatures average from 2 to 4 C.

2.2 Agriculture

Irrigated farms in the study area are served by an extensive system of irrigation supply and drainage canals, which are shown in Figure 2.2. Farms in the MIA have access to off-farm drainage facilities, while farms outside of the MIA generally must retain their drainage on the farm. The drainage from farms in the upstream or eastern half of the MIA (including the Mirrool and Yanco Irrigation Areas and the Benerembah and Tabbita Irrigation Districts) flows into Barren-Box Swamp (see Figure 2.2). After blending with better quality water from the Murrumbidgee River, this drainage water becomes the irrigation supply for the downstream or western half of the MIA (Wah Wah Irrigation District).

There are generally two types of farms in the study area: mixed farms, which are larger and produce both field crops and livestock; and horticulture and vegetable farms, which are smaller and produce higher value tree crops, vine crops and vegetables. Agricultural land use for a typical year is shown in Table 2.1. The major mixed-farms crops are rice, wheat and pasture, which are usually grown in rotation. Horticulture crops such as wine grapes and citrus and other fruit trees are perennial and may take several years before reaching maximum production potential.

The traditional rotation sequence in the study area is a rice-wheat-pasture rotation with a period of about six years. For example, on a particular field a six-year rotation cycle would start with growing rice for two summers (one crop per year), then wheat for two winters and then pasture for two years. Then the cycle would be repeated. There is considerable variability in rotation schedules among different farms and even within a farm on different fields. The benefits of crop rotation are improved weed and disease control, and reduced need for fertilizers. Legume crops and pastures build up residual nitrogen in the soil that reduces the need for nitrogen fertilizer in the next phase of the rotation sequence.

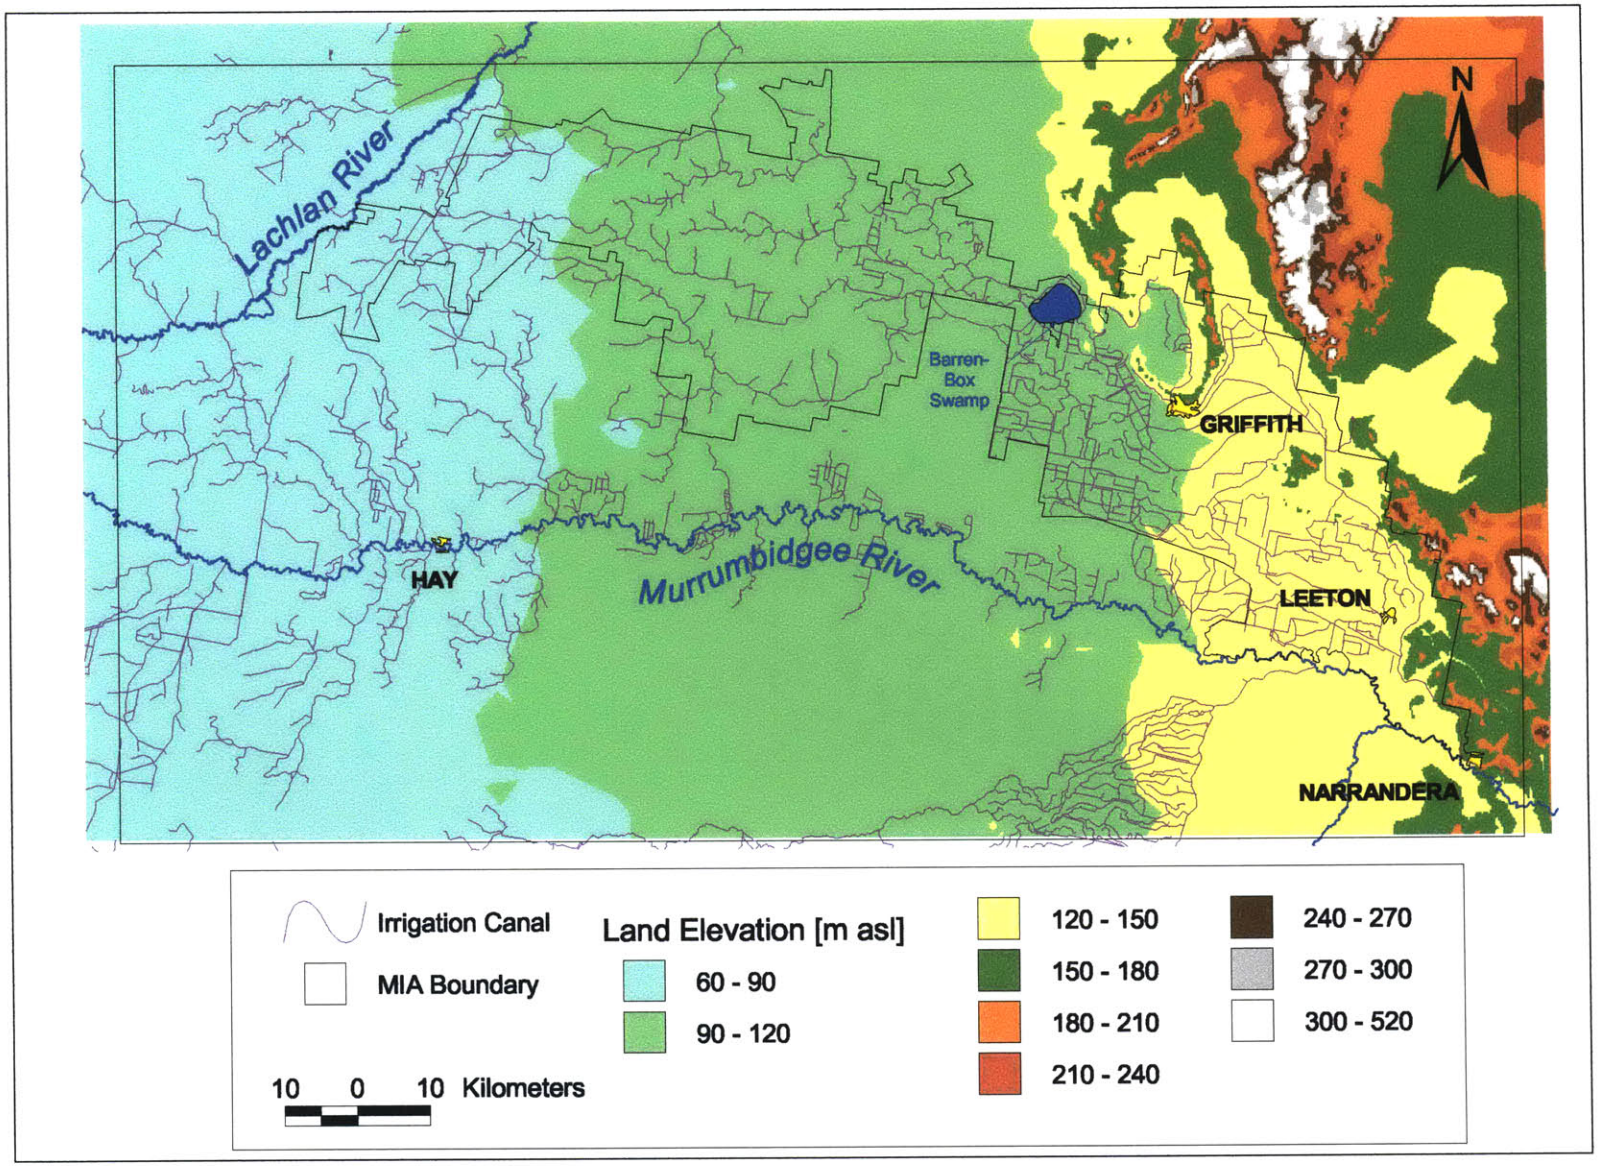

Figure 2.2: Study Area Surface Topography and Irrigation Canals

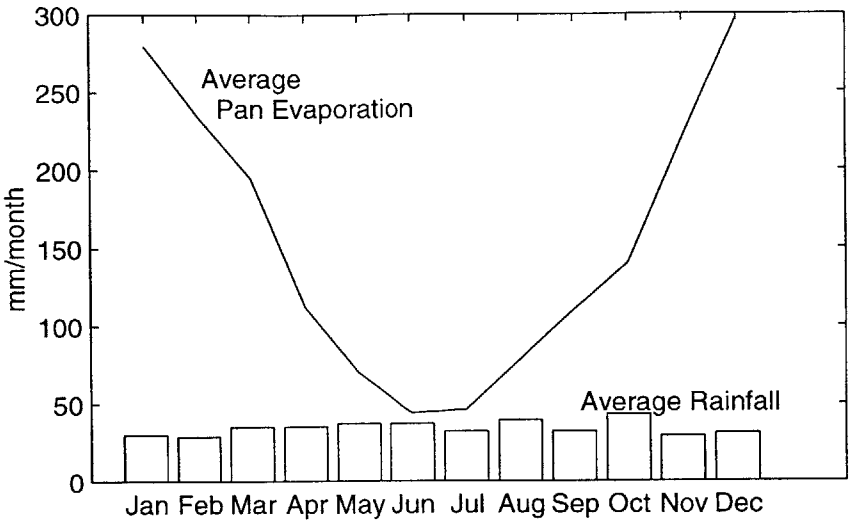

Figure 2.3: Average Monthly Rainfall and Class A Pan Evaporation
Source: Australian Bureau of Meterology [1988]

Although this study focuses on irrigated agriculture, there are extensive areas of dryland wheat and pasture production. In all of the irrigation areas and districts in the Lower Murrumbidgee, dryland production takes place on approximately 309,000 ha. Outside of the irrigation areas and districts, there is about 190,000 ha devoted to dryland cropping [Hall et al., 1993]. As shown on Table 2.1, within the MIA there is at least 78,600 ha of non-irrigated cropland.

The MIA comprises approximately 700 mixed farms with an average size of 250 ha, 1000 horticulture farms with an average size of 20 ha and 50 vegetable farms with an average size of 100 ha [NSW Department of Land and Water Conservation, 1998a]. Farm size in the irrigation areas was originally regulated by the government. These regulations are no longer in force, but they are a major reason that farm size is small in the area.

2.3 Soils

At the eastern edge of the study area, the Murrumbidgee River leaves the foothills at Narrandera and flows out to the Riverine Plain, a vast plain covering much of south-eastern Australia. The Riverine Plain is a relic landform from a previous geologic period when the streams were subject to extensive flooding. During these times the Murrumbidgee broke into many distributaries which spread over the study area in all directions [Langford-Smith and Rutherford, 1966]. Flooding from these prior stream channels gradually built up a vast plain of riverine deposits which merged with the former flood plains of the Murray River to the south and the Lachlan River to the north. The present Murrumbidgee has only one distributary, Yanco Creek, which only flows regularly because of the weir near Leeton. The present floodplains are restricted to the area of the Lowbidgee Irrigation district, considerably west of the study area (see Figure 2.1).

Table 2.1: Typical Land Use in the Study Area by Region

Farm Type and Crop	Upstream MIA[a] [ha]	Downstream MIA[b] [ha]	Outside MIA[c] [ha]
Irrigated Mixed Farming			
Rice	36,200	3,000	10,000[d]
Wheat/Winter Grains	35,000	2,500	0
Annual Pasture	40,000	15,000	35,000
Soybeans/Summer Grains	2,800	2,000	0
Lucerne/Perennial Pasture	4,000	2,500	0
Fallow/Not Currently Irrigated	13,500	3,000	5,000
Total Land Laid Out for Irrigation	131,000	28,000	50,000
Non-Irrigated Mixed Farming			
All Non-Irrigated Crops and Pasture	41,000	37,600[e]	95,000
Total Mixed-Farm Cropland	172,000	65,600	145,000
Horticultural and Vegetable Farming			
Citrus	6,100	0	0
Wine Grapes	9,600	0	0
Stone Fruit	300	0	0
Vegetables	5,000	0	300
Total Horticulture/Vegetable Cropland	21,000	0	300

[a] Includes Yanco, Mirrool, Benerembah and Tabbita Irrigation Areas and Districts. Data adapted from NSW Department of Land and Water Conservation [1998a].

[b] Includes Wah Wah Irrigation District. Data adapted from NSW Department of Land and Water Conservation [1998b].

[c] Includes the part of the study area not in a formal irrigation area or district. Data estimated from Hall et al. [1993] assuming Outside MIA region accounts for 50% of total Lower Murrumbidgee production outside of irrigation areas and districts.

[d] Estimated from aerial photographs.

[e] Does not include dryland cropland in the western part of Wah Wah Irrigation District.

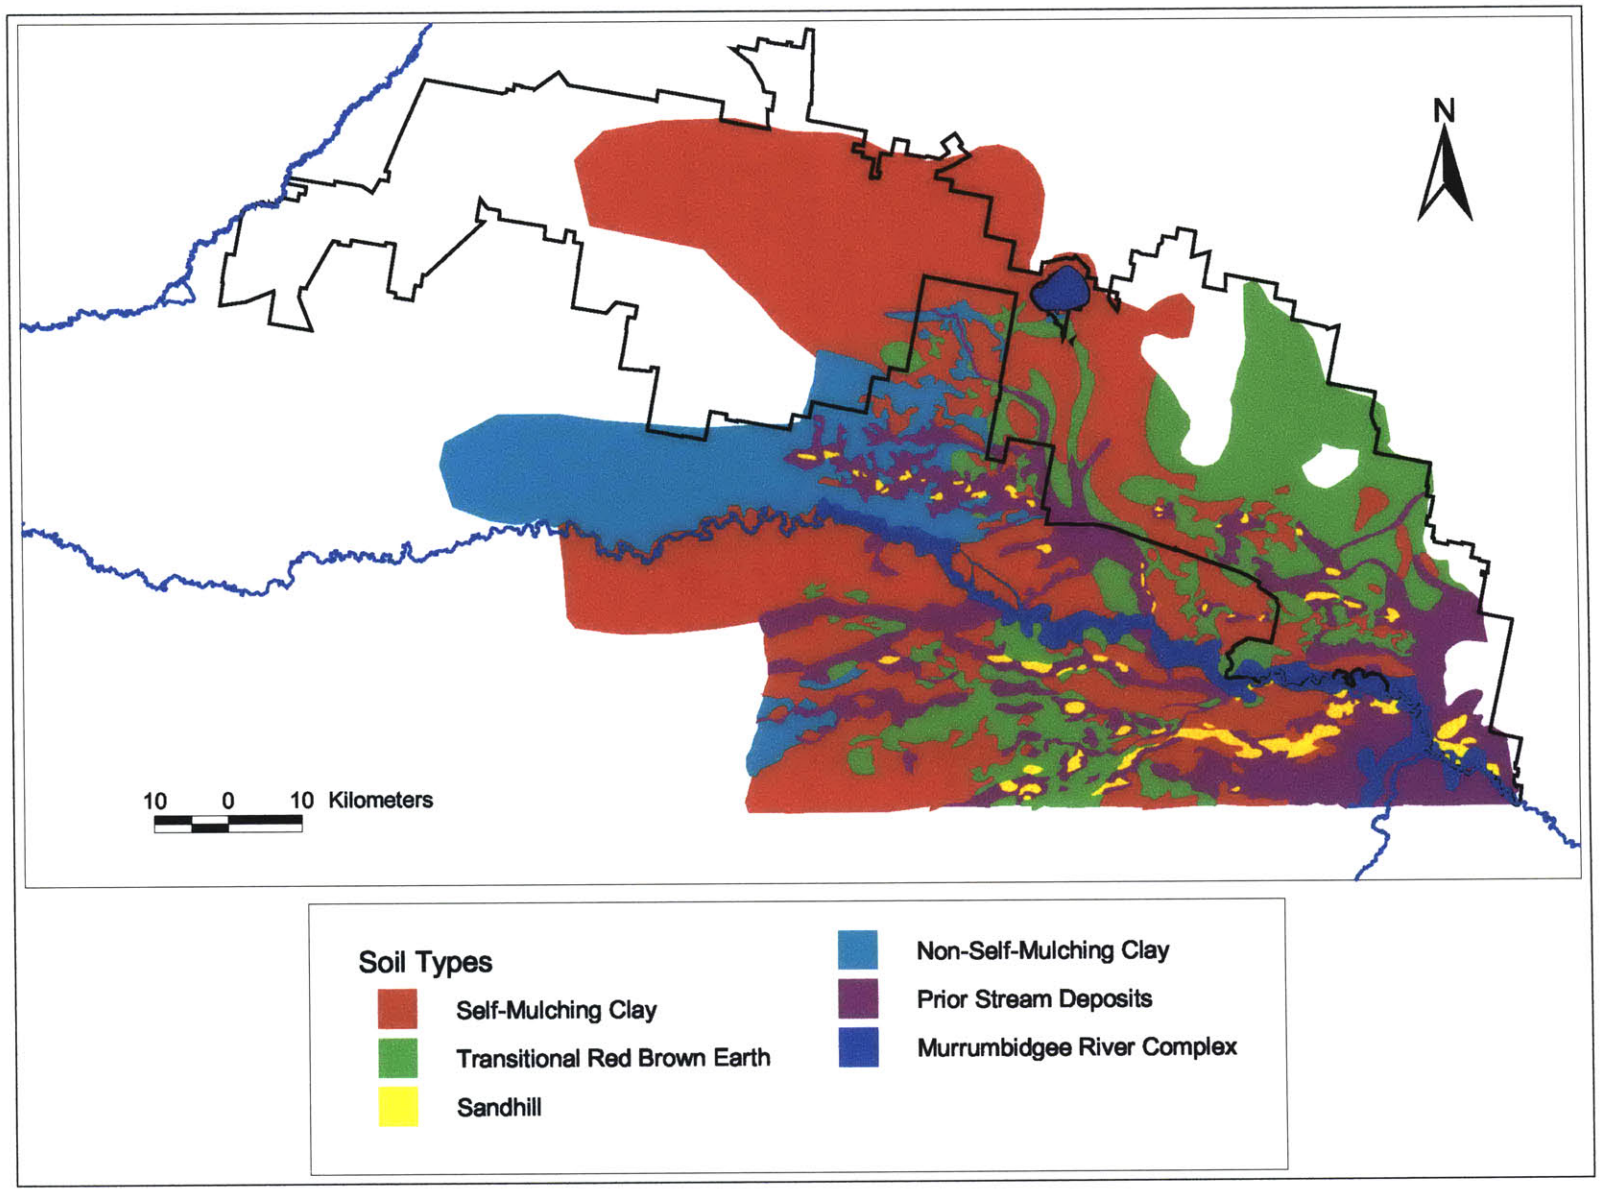

Figure 2.4: Soil Types

The soil types of the mixed-farming areas are shown in Figure 2.4. General soil type data was digitized from a map showing physiographic units of the eastern half of the study area [Stannard, 1966]. No comparable soil data were available for the western half, but since there is little intensive irrigation in this area the soil data were not essential.

The locations of the prior streams are not readily apparent today. The area appears quite uniformly flat with homogeneous clayey soils. There are, however, different soil types which are largely determined by the pattern of prior stream systems. Prior stream deposits are characterized by sand and gravel laid down by the more energetic prior streams. Near the prior streams the soils are often clayey near the surface, becoming more sandy at depth. Sand dune systems are also associated with prior streams.

The soils in the higher plains and low-gradient slopes adjoining the foothills are called transitional red-brown earths. These are duplex soils with a clay-loam A-horizon of around 5-10 cm and a clay B-horizon of low permeability [Olsson and Rose, 1978]. The lower flood plains soils are self-mulching and non-self-mulching clays. These soils are characterized by a heavy texture and a uniform clay profile. Self-mulching means that when the soil surface is dry, extensive cracking occurs breaking the soil into small aggregates. This self-mulching property is good for seedling establishment.

These flood-plain soils generally have poor physical characteristics, low infiltration and internal drainage, and a tendency for the subsoil to disperse. Low soil permeability prevents irrigation water from penetrating very deeply. Since the soil has little capacity to store water, more frequent irrigation is necessary. A more serious problem is that drainage of the profile is so slow that the rootzone remains waterlogged after rain or irrigation. The extent to which waterlogging affects crop production depends on the permeability of the subsoil, the slope of the land, root distribution and the crop tolerance to waterlogging.

The type of crop that can be grown is constrained by the soil type. The most successful agricultural enterprises on the heaviest soils have been rice and pastures which are not as adversely affected by poor soil properties. The locations of the rice fields are shown in Figure 2.5. The areas with self-mulching clays (around 20% of the MIA) are suitable for a wider range of crops, including vegetables and soybeans. The deeper, lighter soils of the prior stream beds and levees are used for citrus, lucerne and some vegetables which need deep soil and good drainage.

Along the hill slopes colluvial soils are commonly found. These soils tend to be more permeable and have better physical characteristics due to the presence of an aeolian clay. The more permeable of these soils near Griffith and Leeton have been developed for irrigated horticulture. The location of horticultural areas is shown in Figure 2.5. Horticultural crops require good drainage, so that almost all of the horticulture farms use deep subsuface drains to protect the crops against waterlogging and salinization. The drains are typically perforated or slotted plastic pipe installed at a depth of 1.6 to 2 meters with a horizontal spacing of 20 to 30 meters [Muirhead et al., 1996].

It is not feasible to use subsurface drainage on the heavier soils. On these soils, surface drainage is used to reduce the amount of time water is ponded on the soil surface which results in less deep percolation. The fields can be laser leveled to make surface drainage more effective. Another possibility is to plant the crops on raised beds. This essentially

provides shallow surface drains in the furrows between the beds.

2.4 Water Resources

2.4.1 Surface Water

The Murrumbidgee is a regulated river for most of its length. Upstream of the study area, water is supplied from two major storage dams: Burrinjuck, with a capacity of 1,026 Gigaliters (GL), and Blowering, with a capacity of 1,628 GL. Water can also be delivered from storages that are part of the extensive Snowy Mountains Scheme. The average annual flow of the Murrumbidgee River entering the study area at Narrandera is 4,000 GL/year and the flow leaving the study area at Hay is 2,300 GL/year [Jolly et al., 1997]. Diversions from within the study area average 2,200 GL per year, of which 98% is used for irrigation. Some of these diversions supply the Coleambally Irrigation Area and individual farms outside of the study area which are not included in our analysis.

About 1000 GL of water per year is diverted from the Murrumbidgee River at Berembed and Gogeldrie weirs. These diversions supply the portion of the MIA upstream of Barren Box Swamp (Yanco, Mirrool, Tabbita and Benerembah) as well as the urban centers of Griffith and Leeton. Except for a portion of the Yanco area which drains back to the Murrumbidgee, returns flows drain to Mirrool Creek and then to Barren Box Swamp (BBS). Benerembah Irrigation District partially reuses the drainage flows before they reach Barren Box Swamp.

As mentioned previously, the drainage water stored in the BBS provides a majority of the supply for the downstream MIA (Wah Wah). BBS now covers an area of 2,800 ha and is permanently filled with water impounded by a levy. Prior to the development of irrigation, it was a natural depression which held water only when local runoff caused Mirrool Creek to flow. Beyond Barren Box Swamp, Lower Mirrool Creek is a series of ephemeral streams and depressions which only reach the Lachlan River during extreme floods.

Approximately 165 GL of the water diverted is used outside of the MIA but within the study area. This water is pumped from the river into private irrigation supply channels. The remainder of the diverted water is used outside of the study area in the Coleambally Irrigation Area and by private irrigators.

2.4.2 Groundwater

Hydrogeology

The study area is in the Riverine Province of the Murray geologic basin. The Murray Basin is a low-lying, saucer-shaped basin of about 300,000 km^2 covering much of southeastern Australia. It is a closed groundwater basin, containing 200-600 m of Cenozoic unconsolidated sediments and sedimentary rocks which form a number of aquifer systems. For the most part, groundwater is trapped in the basin and can only discharge to the surface where it is removed by evaporation and leakage into the river system [Brown, 1989].

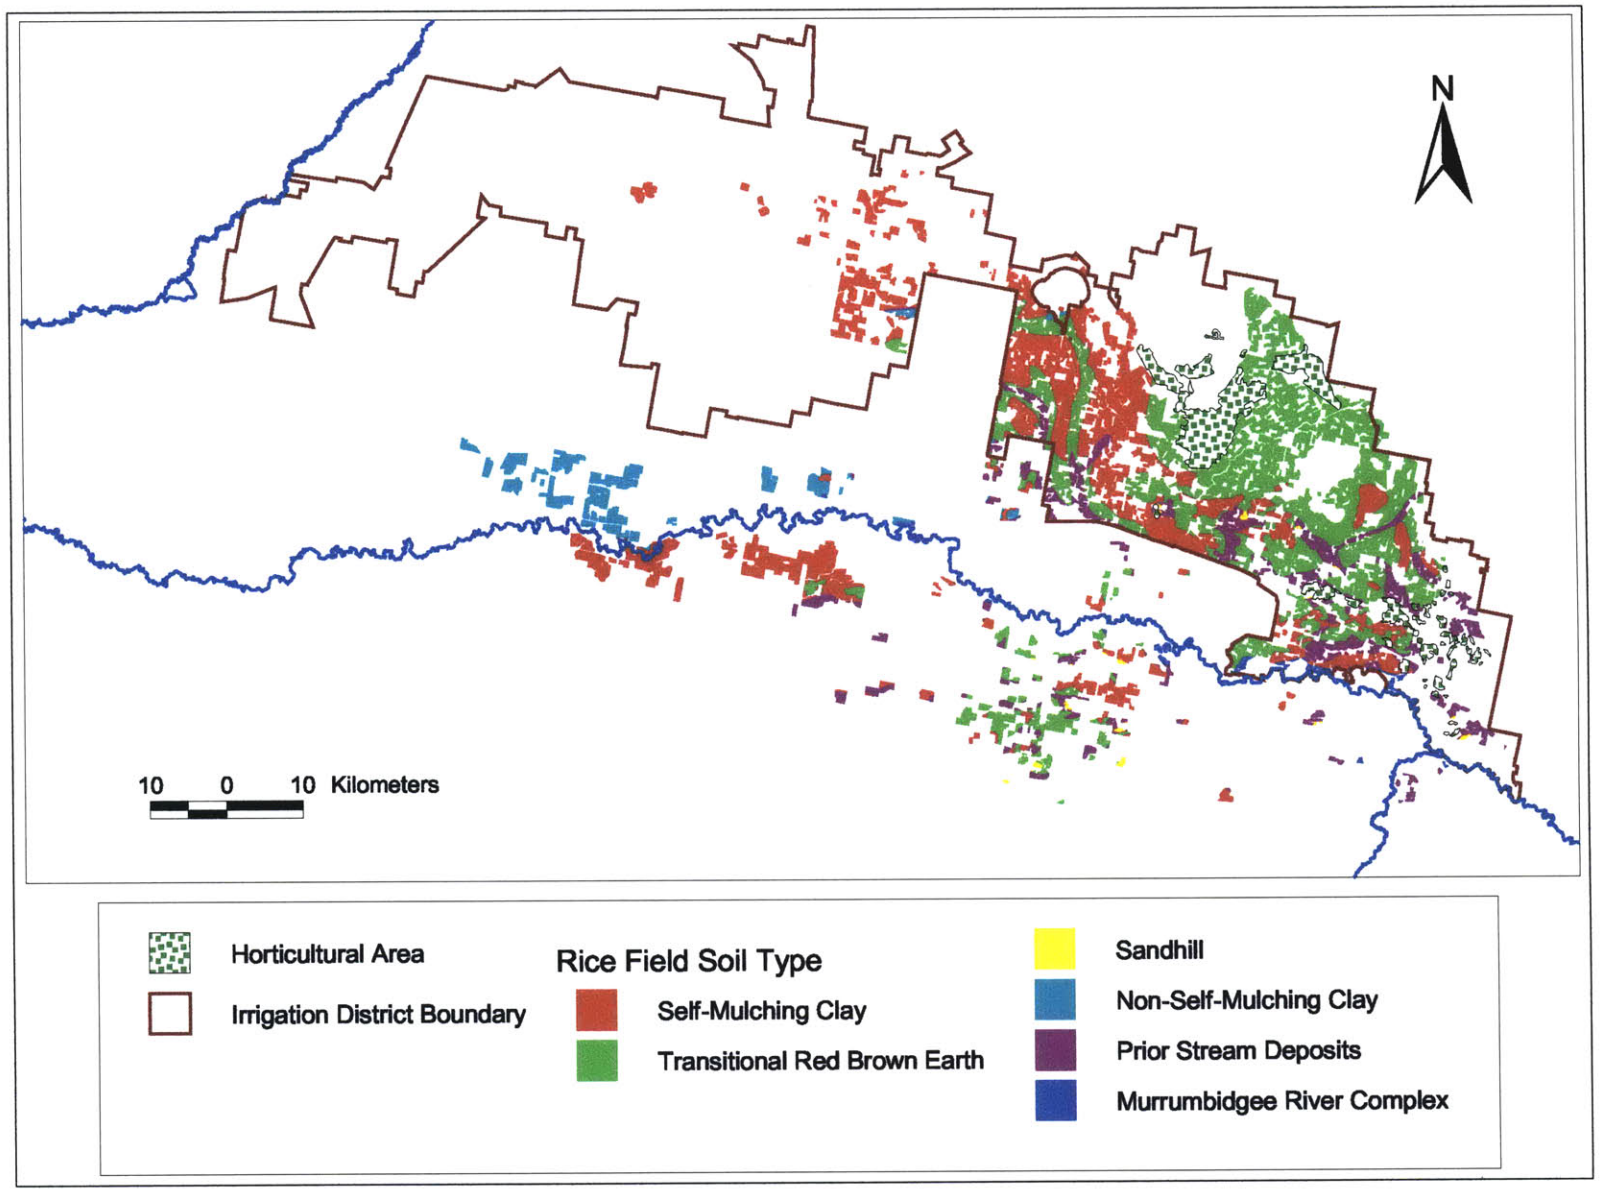

Figure 2.5: Soils Type and Rice Fields

The study area is underlain by unconsolidated alluvium up to 400 meters thick. There are a number of regional aquifer systems. A large fan-shaped area of 6,500 km^2 extending about 120 kilometers downstream of Narrandera is underlain by thick, high-yielding sand and gravel formations, with interbedded silt, clay, peat and brown coal. Groundwater flows from east to west down valley under gentle gradients. Estimated rates of flow in the deep aquifers are 2 to 3 cm/day [Lawson and Webb, 1998]. A more detailed description of the hydrogeology of the study area is given in Section 4.2.

Groundwater Usage

The pattern of groundwater usage in the study area is largely determined by groundwater salinity and potential yields. Within the irrigation areas, groundwater use is minimal due to the availability of lower cost surface water. The best quality groundwater and highest yielding aquifers are in the Murrumbidgee alluvial fan area in the eastern part of the study area near the Murrumbidgee River. This is where most of extraction is taking place, mainly from deeper aquifers which start from 50 to 70 m below the ground surface. Shallower aquifers generally have lower yields and higher salinities. Groundwater usage has been increasing in recent years, as shown on Figure 2.6.

Groundwater Regulation

Groundwater use for irrigation is regulated through a well licensing system. Users are granted an entitlement to pump a nominal volume of water per year. The amount the user is allowed to pump in a given year depends on an announced allocation which is set by the NSW Department of Land and Water Conservation. If the announced allocation is 100%, then the user can pump the amount of their entitlement. From 1991 to 1998, an announced allocation of 150% was in place. Since July 1998, the allocation has been 100%. Figure 2.6 shows the amounts allocated assuming a 100% announced allocation.

A groundwater management plan for the Lower Murrumbidgee Groundwater Management Area is being prepared. This Management Area is approximately the same as the area shown in Figure 2.1. The groundwater system was identified as being at high risk for over-allocation, well interference and the transport of saline groundwater into regions of good quality groundwater. The current level of use may be close to average recharge rate. While the plan is prepared, there is a moratorium on the issuing of new pumping allocations beyond the current total of 494,000 ML [Lawson and Webb, 1998].

2.5 Salinization

2.5.1 Watertable Depths

Groundwater levels were greater than 20 meters below the land surface when irrigation began. Clearing of the native vegetation and large-scale irrigation led to substantially larger recharge rates, particularly in areas where sandy aquifer systems occur close to the surface.

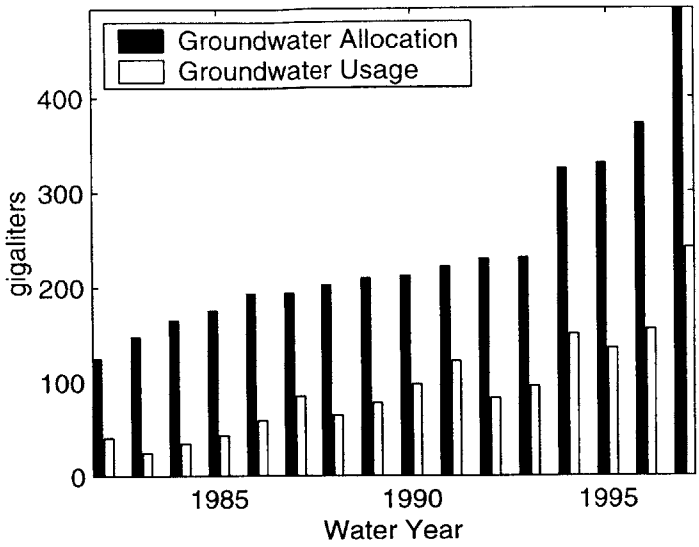

Figure 2.6: Annual Groundwater Usage and Allocation
Source: Lawson and Webb [1998]

Watertables began to rise and threaten agricultural productivity with more frequent waterlogging and salinization. Current areas with high watertables are shown in Figure 2.7. In the MIA, it is estimated that watertables are within 2 m of the surface in over 70% of the area [MIA LWMP Taskforce, 1998]. Areas outside of the formal irrigation areas and districts do not currently have problems with high watertables.

2.5.2 Land and Water Management Plan

There is a great deal of concern among farmers and resource managers in the Murrumbidgee about the economic effects of high watertables. A Land and Water Management Plan for the MIA is currently being negotiated in order to deal with this and other issues. According to the plan, the costs of soil salinization to agriculture will be $26.2 million over the next 30 years [MIA LWMP Taskforce, 1998]. It is estimated in the Plan that 18% of the MIA experiences some degree of crop yield reduction due to salinization. Approximately two and a half percent of the MIA experiences total crop loss. Based on extrapolations of groundwater levels, crop losses are expected to occur on up to 28% of the MIA after 30 years if no action is taken.

The Plan recommends reducing groundwater recharge in order to allow watertables to drop. There are many specific recommendations which include reducing losses from the irrigation delivery system, encouraging better on-farm water management and the adoption of best management practices.

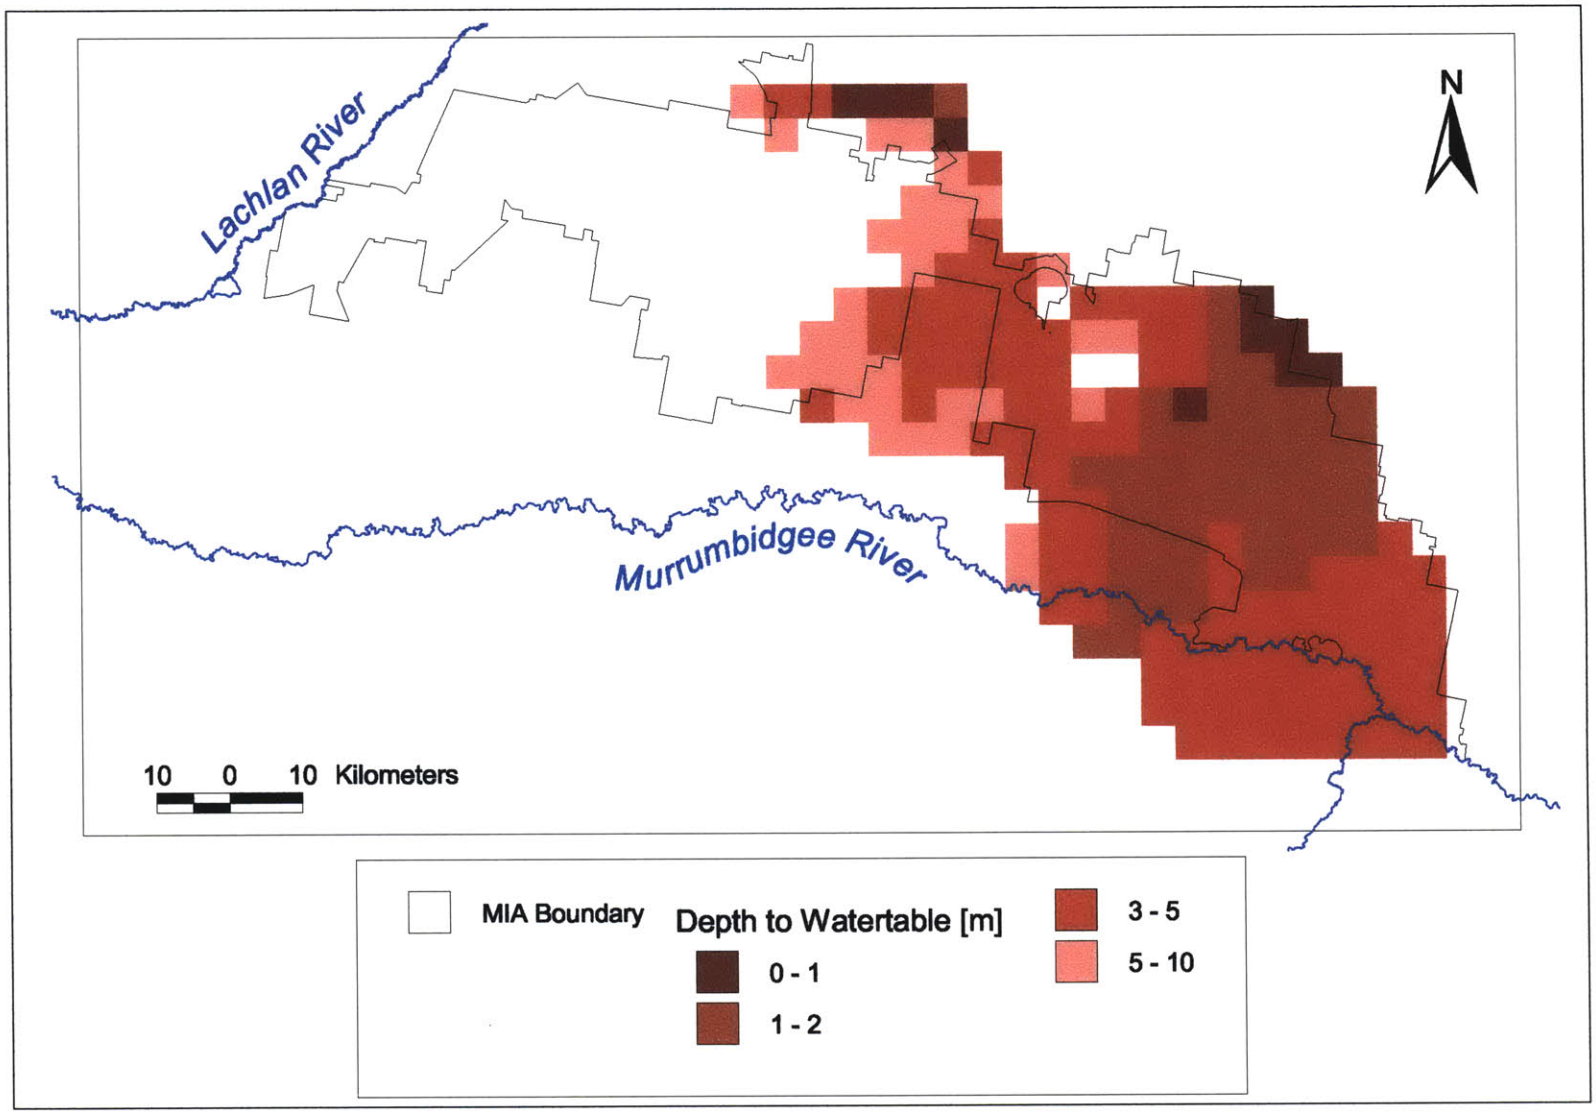

Figure 2.7: Current High Watertable Areas

2.6 Land and Water Management Policies

2.6.1 Recharge Management Under Rice

The New South Wales Department of Land and Water Conservation (DLWC) is responsible for regulating agricultural water and land use. This regulation has focused on rice growing from the very beginning. The history of rice regulation is described in Humphreys et al. [1994]. A brief summary is given in this section.

Rice was first grown in the MIA in 1924, although the MIA was opened in 1912. The construction of Burrinjuck Dam on the Murrumbidgee River in 1928 and Blowering Dam on the Tamut River in 1968 lead to a large increase in irrigation in the Murrumbidgee Valley. A predecessor to DLWC, the Water Conservation and Irrigation Commission, was responsible for rural water projects. It actually sponsored the first commercial rice plantings. Rice areas were restricted by supply channel capacity and the size of the domestic market. From 1933 to 1943, only 33 ha of rice were allowed per farm.

During World War II, the federal government requested additional rice production, which lead to more farmers growing rice, but still under regulation. It was during this time that waterlogging problems developed in parts of the Yanco Irrigation Area, particularly in areas where rice was grown on or near land underlain by shallow aquifers. In response to this, rice growing was restricted to areas where previous water use was less than 27 ML/ha (2.7 m) and the watertable was more than 1.8 m deep. In addition, each farmer could only crop 24 ha of his farm if it was underlain by shallow aquifers and 40 ha if it was not [Muirhead et al., 1990]. These regulations were relaxed in the late 1960s. Total rice area remained at about 15,000 ha from the end of the war until the early 1970s.

A rapid increase in MIA rice area began in mid-1970s when the regulation was changed to allow 73 ha of rice per farm. The change was a result of strong lobbying by the rice industry. By the early 1980s, around 40,000 ha was sown to rice. There have been fluctuations, but the area has remained about at this level since then. In 1996, about 42,000 ha was planted to rice which yielded 361,000 metric tons. The value of this production is AUS$65 million [MIA LWMP Taskforce, 1998].

Starting in 1989, rice growing has been allowed outside of the irrigation areas and districts. However, the environmentally-based restrictions were gradually tightened. The present regulations went into effect in 1993 and are based on the "hydraulic loading" concept. The focus is on the amount of land growing rice as opposed to the irrigation water use or deep percolation from the rice. Each farm in the MIA is allowed to grow rice on 30% of the rice approved land or 65 ha, whichever is greater. For farms outside of the MIA, 100 ha of rice is allowed for every 972 ML of surface water or groundwater that is licensed.

As part of its compliance monitoring of the rice acreage restriction, the NSW Department of Land and Water Conservation takes aerial photographs of the Lower Murrumbidgee Valley every year. The photos are taken in November or December after the rice fields have been flooded, which makes it easier to distinguish rice from other crops. The outlines of the rice fields are digitized and the area of rice within each farm's boundaries is determined using a GIS. Farmers who have grown more than the allowed acreage are fined. Figure 2.5 is based

on the DLWC's work but it includes all rice fields used in the past six years as opposed to a single year [NSW Department of Land and Water Conservation, 1998c].

Land is approved for rice growing based on an analysis of the soil profile based on soil borings. If there is at least 2 m of medium to heavy clay within the top 3 m of the profile, then the land is approved for rice growing. All of the flood plain soils (transitional red-brown earths, self-mulching and non-self-mulching clays) can be approved for rice. There is also a guideline that rice only be grown on fields which use less than 16 ML/ha for a rice crop. This is based on the assumption that 12 ML/ha is consumed in evapotranspiration and the remaining 4 ML/ha becomes surface drainage and deep percolation. Only 100 ha in the MIA has been prohibited from growing rice based on the water use guideline.

2.6.2 Water Allocation System and Water Trading

Prior to the recent introduction of water trading, the water available to irrigators was determined by an administrative system of water allocations. The allocations were based on the nominal quantities of water that are specified on the licenses held by irrigators. During each irrigation season, the Department of Land and Water Conservation made judgments about the amount of water available, and formally allocated this to irrigators. Some irrigators were consistently allocated as much water as they were likely to use, while others generally received the balance of available water on a pro-rated basis.

As demand for water in the Murrumbidgee and other basins increased, the need for a more efficient and flexible approach to allocating water resources became apparent. Studies of the impact of making rights to water in the basin tradable indicated annual net benefits to irrigators of around $50 million for the Murray-Darling Basin as a whole [Hall et al., 1993]. The results of these studies suggested that water management agencies did not always allocate water resources to the irrigators or regions that valued them highest. In general, holders of rights to water in the Murray Darling Basin are now able to trade water on both a permanent and temporary basis, although various constraints and limits on transfers apply.

In the Murrumbidgee region, rights to water are granted in two main forms: high security licenses and general security licenses. At present, holders of high security licenses are guaranteed 100 percent of their licensed water right or entitlement in all but the most serious drought events. The percentage allocation offered to general security license holders is based on the amount of water available after high security licenses (and the needs of other priority water users) are taken into account. In the past, the proportion of entitlement offered to general security license holders has varied from well above 100 percent, to less than 100 percent. With growing demand for water in the region (including the demand for water for environmental and recreational uses), along with significant amounts of unused water licenses, future allocations are expected to remain below 100 percent. Additional "off-allocation" water is sometimes also available. This is extra water from flood flows and dam releases which can be purchased when it is available but that is not counted as part of the allocation.

Starting in 1994, temporary trading of surface water allocations has been significant. This trade has been driven by a combination of factors: 1) the imposition of a cap in 1996 on new

diversions from the Murray-Darling River Basin, 2) the reduction of "off-allocation" water due to a series of drought years, and 3) the marketing of allocations which were previously rarely or never used (known as sleeper and dozer licenses). In the 1997 season, 68,000 ML was traded within the MIA and 98,000 ML was traded from the MIA to other areas including a small volume of water to areas outside of the catchment [NSW DLWC, 1999]. Permanent trade of surface water and both permanent and temporary trade of groundwater has been very limited due to uncertainty about future government policy.

2.6.3 Salinization Policy Issues

Current policies to reduce the impacts from soil salinity focus on regulation of the land planted to rice and the amount of water used by rice. These policies were established many years ago and may not be appropriate given the current situation. In addition, there has no been integrated analysis of the interactions between the economic and hydrologic systems.

Both the rice-area restrictions and water trading can affect the crop distribution in the study area. This crop mix largely determines the amount of deep percolation which will end up recharging the groundwater system. If there is more recharge in the future in high watertable areas, then the salinization problem may become worse. We will analyze the effects of these two policies on the long-term crop choices and watertable depths using the hydrologic-economic model formulated in the next chapter.

Chapter 3

Hydrologic-Economic Model

3.1 Introduction

This chapter begins with a discussion of the how externalities and common property resource use is modeled in an economic framework. We then present a general formulation of our hydrologic-economic model of irrigated agriculture. The model is formulated as a dynamic optimization or optimal control problem in which regional net benefits of agriculture are maximized subject to a set of hydrologic and economic constraints. The main hydrologic constraint is a simple state equation which represents the response of the groundwater system to agricultural production. The main economic constraints are production functions which give total crop yield and deep percolation as a function of shallow groundwater conditions.

3.2 Economic Modeling of Externalities

3.2.1 What is an Externality?

An externality is a form of market failure that occurs because there are no markets for many environmental resources or pollution. These effects, which can be either positive or negative, are *external* to the market system. In order for a system of competitive markets to achieve a (Pareto) efficient allocation, there can be no externalities. A market failure means that resources may not allocated efficiently and some form of intervention to correct the misallocation may be beneficial. However, the existence of externalities does not necessarily lead to inefficiency. The parties involved may privately negotiate a solution without any involvement by the government. In cases where an externality leads to inefficient resource allocation (e.g., too much pollution), it may be in society's interests to try to correct the misallocation.

Policies that address externality problems can be categorized into two groups: command-and-control instruments and economic-incentive instruments. Command-and-control regulations attempt to directly dictate the production process. This is the most commonly used approach and it includes the setting of technology and performance standards. Economic-incentive, or market-based, instruments work indirectly through the influence of market

signals on farmer behavior. Examples include tradeable permits and pollution charges. This category also includes policies which encourage the creation of missing markets through the assignment or reassignment of property rights. Resource use and pollution generation are not directly under the control of the government as is the case with command-and-control instruments. This means that if the policy-making agency does not exactly know the costs and benefits functions faced by the farmers, the result of the policy will be uncertain. Economists have extensively analyzed the theoretical merits and practical limitations of these different policies [e.g. Weitzman, 1974; Baumol and Oates, 1988; Hahn and Stavins, 1992].

In this study we are concerned with production or technological externalities, specifically, the effect of high water tables on crop production. The drainage water that farmers discharge to the groundwater system is not priced by any market. But the drainage water of one farmer contributes to the high watertables in the region, which affects many other farmers. Note that this externality is not always negative. Higher watertables may lead to lower irrigation water requirements, which means lower production costs. Only farmers in the area can discharge drainage to the groundwater system, but there is no mechanism for these farmers to compensate each other for the effects on other farmers. Therefore, we can consider the drainage capacity of the groundwater system as a common pool resource.

Externalities are also classified as stock or flow externalities. The effect of high watertables is an stock or intertemporal externality because the external damages occur in a future period once the resource stock (watertable depth) has reached a low enough level. A flow externality, in contrast, causes external damages immediately. We do not consider any flow externalities in our analysis, but this is the type considered in most economic studies of externalities. Stock externalities are much more difficult to model because we must represent the dynamics of resource use. The fact that groundwater recharge from one farm affects watertables on all other farms in the future is why we have to use a dynamic and spatially-distributed framework.

3.2.2 Quantifying the Cost of Externalities

The inefficiencies caused by externalities lead to a reduction in social welfare. In our case, we focus our attention on agricultural production in the study area; therefore, our measure of social welfare is the total return to irrigated agriculture in the region. There are many other potential contributions to social welfare which we do not consider.

A socially optimal resource allocation is the one that maximizes social welfare. We can quantify the effect of the externality by estimating the loss of total welfare resulting from the inefficiencies. If this loss is large compared to the costs of correcting the externality, then intervention is justified on economic grounds. Our goal in this study is to estimate the potential benefits to addressing the externality problem. We do not attempt to estimate the costs of potential interventions, although these must also be taken into account when deciding on a policy. However, if the potential benefits are small, then intervention will not significantly increase social welfare and may actually decrease it. Our approach to quantifying the effect the drainage externality is to construct a mathematical model of the combined hydrologic-economic system. We use the model to simulate different scenarios and then compare the social welfare estimates of the scenarios.

When using a common-pool resource such as a groundwater aquifer, farmers pay only their own private costs and ignore the costs to others of increased scarcity. The additional cost is referred to by economists as a user cost or a scarcity rent. The end result is that the resource may be used at a higher rate than is socially optimal. The common-pool resource issue can also be considered a stock externality.

3.2.3 Models of Behavior Under Common Property Arrangement

There are two main models of behavior which can be assumed under a common property regime. The first is that the farmers ignore the state of the groundwater system and decide on their production using myopic, open-loop policies. This approach is the most straight-forward to implement and is used in most applied analyses of both groundwater extraction [e.g. Gisser and Sanchez, 1980; Feinerman and Knapp, 1983; Allen and Gisser, 1984; Nieswiodomy, 1985; Worthington et al., 1985] and agricultural recharge [Shah et al., 1995]. The other potentially more realistic approach is to assume that the users of the groundwater resource develop closed-loop policies, which depend on the current state of the system [e.g. Negri, 1989]. This type of model is more difficult to solve numerically, so it is mainly used in theoretical work. Resources under the open-loop models are likely to be depleted faster than with closed-loop models. Therefore, simulating the myopic case provides a worse-case scenario which can be compared to the best-case scenario of socially optimal resource use.

3.3 Hydrologic-Economic Model Formulation

Our mathematical model of the hydrologic-economic system is formulated as a discrete-time optimal control problem. The objective function is the discounted sum of total returns from agriculture. This will be maximized subject to constraints which describe the groundwater system, the unsaturated system and the economic system.

Due to data and resource constraints, we have not included the Coleambally Irrigation Area in the model. Since the characteristics of the CIA are sufficiently similar to the MIA it is likely that expanding the scope of the model to include the CIA would not significantly change the results.

3.3.1 Major Assumptions

We assume that the prices of crops and production inputs are exogenous and unchanged by the policies we are analyzing. This is referred to as a partial equilibrium analysis in economics. We are considering only one sector of the economy, agricultural production, and ignoring any secondary effects outside of this sector. If the production levels change significantly, there will be some secondary effects in sectors which are closely linked to agriculture, such as the agricultural labor or agricultural chemical sectors. Since our results do not predict large changes in the production levels, it is reasonable to neglect secondary effects.

Our analysis is limited to the mixed-farming areas. We do not include the horticultural areas since they are protected from salinization by subsurface drainage systems. We assume that there will be no radical changes in crop production in the region.

The effects of economic and hydrologic uncertainty are neglected in our analysis. Economic uncertainty includes our lack of knowledge about future prices and the rate of technological innovation. Significant improvements in crop yields or changes in crop prices would undoubtedly change the results. Sources of hydrologic uncertainly include the spatial variability of hydrologic parameters and the stochastic nature of river flow and meteorology. Ignoring these uncertainties in our model will tend to predict resource use that is very aggressive with no "hedging" behavior. If uncertainty were accounted for, the resulting model would tend to use resources more conservatively in order to reduce the possibility of large losses.

3.3.2 Economic Units and Hydrologic Cells

The model area is divided into 37 economic units within which agricultural production conditions are similar. Shown in Figure 3.1, the economic units are the scale on which the economic decisions are made. In our model, there is only one decision: the annual crop mix for each economic unit. The control variables u_t^{ij} are the areas planted to crop j in economic unit i during year t. For each time we collect all the control variables into a control vector U_t,

$$U_t = [u_t^{11}\ u_t^{12}\ \ldots\ u_t^{ij}\ \ldots\ u_t^{IJ}]^T, \quad t = 1, \ldots, T, \tag{3.1}$$

where $i = 1, \ldots, I$ is an index over the economic units, $j = 1, \ldots, J$ is an index over the crops, and t is the time index. The control vector has a total length $M = IJ$. The time horizon of the model is divided into annual time periods $t = 1, \ldots, T$. We also define the vector of control variables for each economic unit and time

$$U_t^i = [u_t^{i1}\ u_t^{i2}\ \ldots\ u_t^{ij}\ \ldots\ u_t^{iJ}]^T, \quad t = 1, \ldots, T;\ i = 1, \ldots, I \tag{3.2}$$

The hydrologic conditions at a given time and region are represented by the state variables x_t^k. The scale of the groundwater model is smaller than the economic units so that the groundwater response can be modeled more accurately. The hydrologic cells are the smaller squares shown in Figure 3.1. For each time t, the state of the groundwater system is specified by the state vector X_t,

$$X_t = [x_t^1\ x_t^2\ \ldots\ x_t^k\ \ldots\ x_t^N]^T, \quad t = 0, \ldots, T, \tag{3.3}$$

where $k = 1, \ldots, N$ is the index over states. Note that there are fewer states than hydrologic grid cells because we used a model order reduction technique. This is described in Chapter 4.

3.3.3 Crop Production and the Unsaturated Zone Model

From our simulation experiments of the unsaturated zone, we estimated three nonlinear functions which describe the interaction between crop production and the groundwater system.

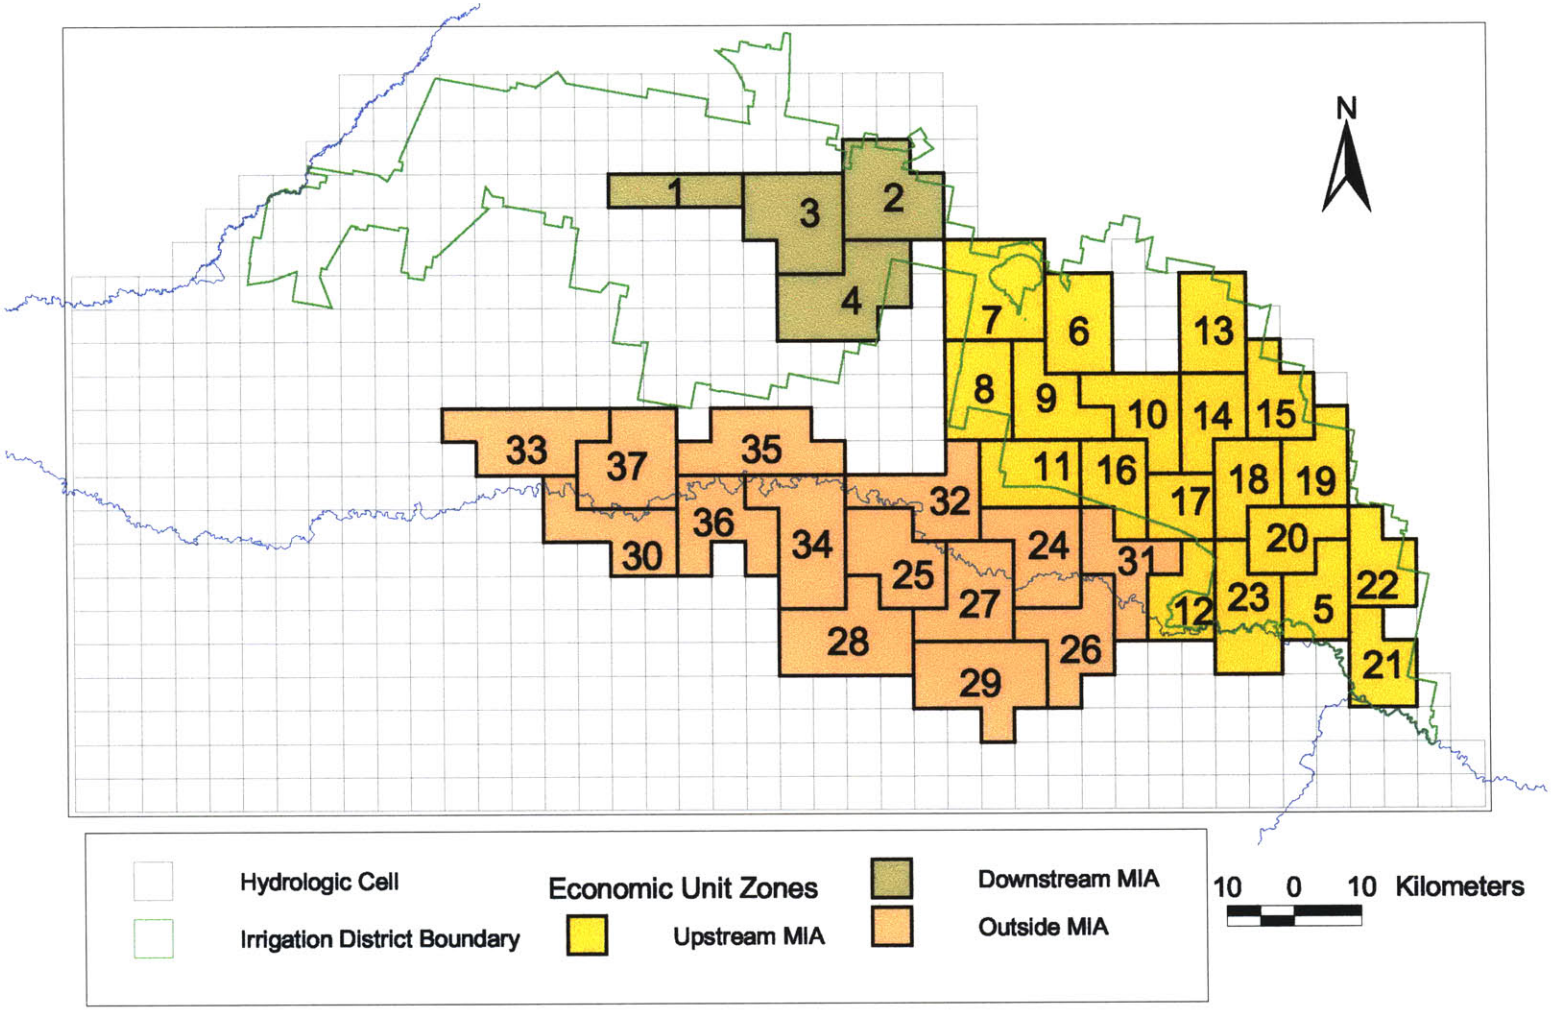

Figure 3.1: Economic Unit Numbers and Zones

The functions give crop yield per unit area y_t^{ij}, groundwater recharge per unit area r_t^{ij}, and irrigation requirement per unit area w_t^{ij} for each crop in each unit at each time. The details are described in Chapter 5, but the form of the equations is

$$y_t^{ij} = f_y(d_t^i, \alpha^{ij}) \tag{3.4}$$

$$r_t^{ij} = f_r(d_t^i, \alpha^{ij}) \tag{3.5}$$

$$w_t^{ij} = f_w(d_t^i, \alpha^{ij}) \tag{3.6}$$

where depth to groundwater d_t^i is the independent variable and α^{ij} is a vector of parameters describing the economic unit and crop characteristics.

The yield function given in Eqn. 3.4 only represents the effects of salinization on crop yields at the field scale with no interactions among crops. There are also other factors which affect yields at the economic unit scale, such as heterogeneous land quality and crop rotations. These factors result in a large-scale yield functions which exhibit decreasing returns to land [Howitt, 1995]. This means that as the land planted to a particular crop increases, the yield per unit area decreases. We represent this effect by multiplying the field-scale yield from Eqn. 3.4 by a factor $\gamma(u_t^{ij})$, which is between zero and one. The large-scale yield is then

$$\bar{y}_t^{ij} = y_t^{ij}\gamma(u_t^{ij}), \qquad \forall\, t, i, j \tag{3.7}$$

We specify $\gamma(u_t^{ij})$ by fitting a polynomial function to observations of yield dependence on planted areas. This is described in Section 5.6.2.

The large-scale yield function for each economic unit can be written as a diagonal matrix

$$Y^i(d_t^i, U_t^i) = \text{diag}[\bar{y}_t^{i1}\ \bar{y}_t^{i2}\ \ldots\ \bar{y}_t^{ij}\ \ldots\ \bar{y}_t^{iJ}], \quad i = 1, \ldots, I \tag{3.8}$$

Similarly, we define a matrix of irrigation requirements for each economic unit

$$W^i(d_t^i) = \text{diag}[w_t^{i1}\ w_t^{i2}\ \ldots\ w_t^{ij}\ \ldots\ w_t^{iJ}], \quad i = 1, \ldots, I \tag{3.9}$$

We also define D_t as the vector of groundwater depths averaged over the economic units at time t,

$$D_t = [d_t^1\ d_t^2\ \ldots\ d_t^i\ \ldots\ d_t^M]^T, \quad t = 1, \ldots, T. \tag{3.10}$$

3.3.4 Groundwater Flow State-Space Model

The form of the groundwater state-space model is

$$X_{t+1} = AX_t + BR(X_t)U_t \tag{3.11}$$

$$D_t = CX_t + \bar{D} \tag{3.12}$$

where A is the $N \times N$ state transition matrix, B is the $N \times I$ input matrix, and C is the $M \times N$ output matrix. The nominal depth to groundwater $\bar{D}$ depends on the nominal

groundwater heads and the ground surface elevation. The $I \times M$ matrix $R(X_t)$ organizes the recharge per unit area values as

$$R(X_t) = \begin{bmatrix} r_t^{11} & r_t^{12} & \dots & r_t^{1J} & 0 & \dots & 0 & \dots & 0 & \dots & 0 \\ 0 & 0 & \dots & 0 & r_t^{21} & \dots & r_t^{2J} & \dots & 0 & \dots & 0 \\ \vdots & \vdots & & \vdots & \vdots & & \vdots & & \vdots & & \vdots \\ 0 & 0 & \dots & 0 & 0 & \dots & 0 & \dots & r_t^{I1} & \dots & r_t^{IJ} \end{bmatrix}. \tag{3.13}$$

We organize $R(X_t)$ in this way so that when it premultiplies U_t the result is an $I \times 1$ vector of the total groundwater recharge for each economic unit at time t. In view of Eqn. 3.12, we use X_t instead of D_t as the argument for R. In the future, we will also write $Y^i(X_t, U_t^i)$ and $W^i(X_t)$ to reflect their dependence on the groundwater states.

3.3.5 Land and Water Resource Constraints

Irrigable Land Constraint

The land cropped in each unit must be non-negative and less than the maximum land available, U^i_{max}. In addition, the area of rice grown must be less than the area of rice-approved land, $U^i_{maxrice}$.

$$u_t^{ij} \geq 0, \qquad \forall t, i, j \tag{3.14}$$

$$\sum_j u_t^{ij} \leq U^i_{max} \qquad \forall t, i \tag{3.15}$$

$$u_t^{i,rice} \leq U^i_{maxrice} \qquad \forall t, i \tag{3.16}$$

We use the rice GIS data shown in Figure 2.5 to estimate where the irrigable land in each economic unit is located. The figure shows the union of the fields where rice was grown over the past six years (1993-1998). We assume that the area of this union is the total amount of rice-approved land. These rice areas also give a good estimate of the location of potential land which can grow other irrigated broadacre crops as well as rice. Most farmers follow a rotation based around rice so this is a reasonable assumption. The amount of irrigable land for the three irrigation zones (Upstream MIA, Downstream MIA, and Outside MIA) is also known. We increase the irrigable area estimates based on the rice GIS data so that it is consistent with the data for the three zones. Table 3.1 shows the data and estimates by zone.

Riceland Restriction Policy Constraint

For model scenarios in which the riceland restriction policy is in effect, the maximum land available for rice, $U^i_{maxrice}$, is reduced to 30% of the rice-approved land.

Table 3.1: Modeled Areas, Irrigable Areas, and Rice Area by Zone

Irrigation Zone[a]	Econ. Unit Area [1000 ha]	Irrigable Area [1000 ha]	Ricefield Area [1000 ha]	Non-rice Irrigable Area [1000 ha]	Number of Econ. Units [-]	Irrigable Area Added per Unit [1000 ha]
Upstream MIA	265.0	131.0[b]	102.3	28.7	19	0.271
Downstream MIA	62.5	28.0[c]	8.6	19.4	4	0.778
Outside MIA	265.0	50.0[d]	32.9	17.1	14	0.161
All Zones	592.5	209.0	143.7	65.3	237	

[a] Irrigation Zones are shown in Figure 3.1.
[b] *Source:* NSW Department of Land and Water Conservation [1998a]
[c] *Source:* NSW Department of Land and Water Conservation [1998b]
[d] Estimate based on Hall et al. [1993].

Water Constraint Without Water Quasi-Market

The form of the water constraint depends on whether there is a water market or not. If the water allocation is fixed to the economic unit and can not be traded to another unit, then the constraint is

$$W^i(X_t)U_t^i \leq W_{max}^i \qquad \forall t, i \tag{3.17}$$

Water Constraint With Water Quasi-Market

If there is a market for water within the region, then there is only one constraint for the entire region

$$\sum_i W^i(X_t)U_t^i \leq W_{max} \qquad \forall t \tag{3.18}$$

where $W_{max} = \sum_i W_{max}^i$.

The availability of water summarized by economic unit zone is shown in Table 3.2. We assume that the current levels of water use will apply over the planning horizon. There is some uncertainty over the future allocations of both surface and groundwater. Surface water allocations to agriculture may be reduced in order to increase river flows for ecological reasons. Groundwater use will probably increase in the future which will at a minimum compensate for reduced surface water supplies. We make the conservative assumption that the total amount of water available will not change significantly.

Resource Constraint Set

These constraints define a convex set $\mathcal{U}_t$ which the optimal controls must belong to. We can summarize these constraints as

$$U_t \in \mathcal{U}_t \subset \mathbb{R}^M \tag{3.19}$$

Table 3.2: Irrigation Water Availability by Source and Zone

Water Source	Upstream MIA [GL]	Downstream MIA [GL]	Outside MIA [GL]
Murrumbidgee River[a]	860	100	275
Deep Groundwater[b]	27	0	146

[a] Source: MIA LWMP Taskforce [1998]; Hall et al. [1993]
[b] Source: Lawson and Webb [1998]

where $\mathcal{U}_t$ is defined by the equations

$$G_t(X_t, U_t) \leq 0. \tag{3.20}$$

3.3.6 Objective Function

The net revenues for each economic unit i at each time t is

$$\pi_t^i(X_t, U_t^i) = [P_c Y^i(X_t, U_t^i) - P_w W^i(X_t) - P_v]U_t^i \tag{3.21}$$

where the first term in brackets is the revenue per unit area from crop production, the second term is the cost per unit area of irrigation water, and the third term is the variable cost per unit area except for irrigation water. The $1 \times J$ vectors P_c, P_w and P_v are the crop prices, water prices and variable costs per unit area. Values for these parameters are given in Tables 3.3 and 3.4. The $J \times J$ diagonal matrices $Y^i(X_t, U_t^i)$ and $W^i(X_t)$ are the previously defined yield per unit area and irrigation requirement per unit area. The dimensions are defined so that π_t^i is a scalar giving the total net revenues for economic unit i at time t.

Table 3.3: Crop Prices and Gross Margins

Crop	Average Yield [ton/ha]	Average Price [$/ton]	Income [$/ha]	Var. Cost W/O Irrig. [$/ha]	Irrigation Usage [ML/ha]	Gross Margin [$/ha]
Rice	8.3	160	1,328	410	15.3[a]	753
Wheat	3.0	123	369	250	4.0	76
Annual Pasture	9.0	17	149	50	5.0	45

[a] Assuming typical irrigation requirement.
Source: NSW Department of Land and Water Conservation [1998a]

The best outcome from society's point of view is to maximize the total revenues over all of the regions. If there were a social planner who could dictate the production in each region, he or she would internalize any externalities and use the following objective function to find the socially optimal outcome

$$J_{opt} = \sum_{t,i} \alpha_t \pi_t^i(X_t, U_t^i) + \alpha_T \Pi(X_T) \tag{3.22}$$

Table 3.4: Water Price By Source and Irrigation Zone

Water Source	Upstream MIA [$/ML]	Downstream MIA [$/ML]	Outside MIA [$/ML]
Murrumbidgee River[a]	11	11	8
Deep Groundwater[b]	11	11	8

[a] Source: NSW Department of Land and Water Conservation [1998a]; Hall et al. [1993]
[b] Estimated.

where $\Pi = \sum_i \Pi^i$ is the total residual value of the stock in the economic units at the final time and α_t is a discount factor. We use the discrete-time version of the standard exponential discount factor

$$\alpha_t = \frac{1}{(1+\rho)^t}, \qquad t = 0, \ldots, T \tag{3.23}$$

where ρ is the discount rate.

Under a common pool regime, the farmers in each economic unit will try to maximize only their own net revenues ignoring the effects of their actions on other economic units. In addition, since the economic units cannot be sure they will be able to enjoy the benefits of a lower watertable in the future, they will simply optimize for the present ignoring the effect of their actions on the future. Essentially, this means they ignore the state equation when optimizing. The farmers may understand groundwater dynamics but it is better for them to ignore this in their decision-making since the depth to groundwater is a common pool resource. Since each economic unit will optimize its own profits over each time, this gives rise to IT coupled optimization problems. The objective functions for each of these problems is

$$J_t^i = \pi_t^i(X_t, U_t^i), \qquad \forall\, i, t \tag{3.24}$$

We are assuming that there is cooperation within economic units. This is equivalent to assuming that the units represent very large farms each having a single owner.

When we want to simulate the socially optimal resource use, we will maximize J_{opt} from Eqn. 3.23 subject to the constraints specified in the previous sections. This will be compared to the solution under the common pool regime found by simultaneously maximizing the objectives of Eqn. 3.24.

3.4 Optimality Conditions

In this section, we examine the optimization problem analytically using optimal control theory to gain some insight into the issues. In Chapter 6 we will examine the numerical solutions.

3.4.1 Social Optimum

The social optimum models the best-case scenario where all externalities have been internalized and the future effects of production are taken into account (intertemporal efficiency). The social planner's problem is

$$\max_{U_t} J_{opt} = \sum_{t,i} \alpha_t \pi_t^i(X_t, U_t^i) + \alpha_T \Pi(X_T) \tag{3.25}$$

subject to the state equation (3.11) and resource constraints (3.20). We begin by deriving the optimality conditions using a Hamiltonian formulation. The current-value (undiscounted) Hamiltonian for this problem is

$$\mathcal{H}_t(X_t, U_t, \Lambda_{t+1}) = \sum_i \pi_t^i(X_t, U_t) + \Lambda_{t+1}^T [AX_t + BR(X_t)U_t] \tag{3.26}$$

where $\Lambda_{t+1}^T = [\lambda_t^1 \ \ldots \ \lambda_t^k \ \ldots \ \lambda_t^N]$ is the vector of current-value adjoint variables.

From the maximum principle, the first-order necessary conditions for an optimal solution are

$$\hat{U}_t = \arg\max_{U_t \in \mathcal{U}_t} \mathcal{H}_t(\hat{X}_t, U_t, \hat{\Lambda}_{t+1}), \qquad t = 0, \ldots, T-1 \tag{3.27}$$

$$\hat{\Lambda}_t = \nabla_{X_t} \mathcal{H}_t(\hat{X}_t, \hat{U}_t, \hat{\Lambda}_{t+1}), \qquad t = 1, \ldots, T-1 \tag{3.28}$$

$$\hat{\Lambda}_T = \nabla_{X_T} \Pi(\hat{X}_T) \tag{3.29}$$

$$\hat{X}_{t+1} = A\hat{X}_t + BR(\hat{X}_t)\hat{U}_t, \qquad t = 0, \ldots, T-1 \tag{3.30}$$

$$\hat{X}_0\text{: given} \tag{3.31}$$

where $\hat{X}_t$, $\hat{U}_t$ and $\hat{\Lambda}_{t+1}$ are the state, control and adjoint variables on the optimal trajectory [Bertsekas, 1999]. This system of equations defines a discrete-time, two-point boundary-value problem. Eqn. 3.28 is called the adjoint equation and its boundary condition Eqn. 3.29 is called the transversality condition. The state equation (3.30) moves forward in time from its initial condition (3.31) while the adjoint equation moves backwards in time from its transversality condition.

The maximum principle can be given an economic interpretation [Dorfman, 1969]. The first term of $\mathcal{H}_t$ is the profit resulting from the crop mix U_t. In the second term, $[AX_t + BR(X_t)U_t] = X_{t+1}$ is future state resulting from crop mix U_t. The adjoint variable Λ_{t+1} can be interpreted as the shadow price of the resource stock or state. When this shadow price is multiplied by the state it gives the value of the remaining resource stock at that time. Thus when we examine the Hamiltonian for different crop mixes, the first term represents the *current-profit effect* of U_t and the second term represents the *future-profit effect* of U_t. By choosing the U_t which maximizes the Hamiltonian at each time, we can find the crop mix which gives us the highest total profits over the planning horizon. The transversality

condition gives the shadow price of the resource at the final time. This can be derived from the salvage value or total price the resource could be sold for at time T.

In order to incorporate the constraints on the controls, we augment the Hamiltonian with Eqn. 3.20 and apply the Kuhn-Tucker conditions. The augmented Hamiltonian is

$$\tilde{\mathcal{H}}_t(X_t, U_t, \Lambda_{t+1}, \theta_t) = \mathcal{H}_t + \theta_t^T G_t(X_t, U_t) \tag{3.32}$$

where θ_t is a vector of Lagrange multipliers for the control constraints at time t. The complete first-order optimality conditions using the augmented Hamiltonian are

$$\nabla_{U_t} \tilde{\mathcal{H}}_t(\hat{X}_t, \hat{U}_t, \hat{\Lambda}_{t+1}, \hat{\theta}_t) = 0 \tag{3.33}$$

$$\hat{\Lambda}_t = \nabla_{X_t} \tilde{\mathcal{H}}_t(\hat{X}_t, \hat{U}_t, \hat{\Lambda}_{t+1}, \hat{\theta}_t) \tag{3.34}$$

$$\hat{\Lambda}_T = \nabla_{X_T} \Pi(\hat{X}_T) \tag{3.35}$$

$$\hat{X}_{t+1} = A\hat{X}_t + BR(\hat{X}_t)\hat{U}_t \tag{3.36}$$

$$\hat{X}_0\text{: given} \tag{3.37}$$

$$G_t(\hat{X}_t, \hat{U}_t) \leq 0 \tag{3.38}$$

$$\hat{\theta}_t \geq 0, \quad \hat{\theta}_t^T G_t(\hat{X}_t, \hat{U}_t) = 0 \tag{3.39}$$

where $t = 0, \ldots, T-1$.

We will examine the first-order condition for $\nabla_{U_t} \tilde{\mathcal{H}}_t$ (Eqn. 3.33) in more detail. First, we substitute in the expression for the augmented Hamiltonian given in Eqn. 3.32. With this substitution and dropping the hats for clarity, we have

$$\nabla_{U_t} \tilde{\mathcal{H}}_t = \nabla_{U_t} \sum_i \pi_t^i + \Lambda_{t+1}^T \nabla_{U_t} BR(X_t) U_t + \theta_t^T \nabla_{U_t} G_t(X_t, U_t) = 0, \qquad \forall t \tag{3.40}$$

If we break the gradient operator ∇_{U_t} into separate components for each economic unit i, we have

$$\begin{aligned} \nabla_{U_t^i} \tilde{\mathcal{H}}_t &= \nabla_{U_t^i} \sum_i \pi_t^i + \Lambda_{t+1}^T \nabla_{U_t^i} \sum_i B^i R^i(X_t) U_t^i + \nabla_{U_t^i} \sum_i (\theta_t^i)^T G_t^i(X_t, U_t^i) \\ &= \nabla_{U_t^i} \pi_t^i + \Lambda_{t+1}^T B^i R^i(X_t) + (\theta_t^i)^T \nabla_{U_t^i} G_t^i(X_t, U_t^i) \\ &= 0 \end{aligned} \tag{3.41}$$

where $i = 1, \ldots, I$ and $t = 0, \ldots, T-1$. We use B^i to refer to the ith column of the $N \times I$ matrix B, and $R^i(X_t)$ to refer to the ith row and J nonzero columns of the $I \times M$ matrix $R(X_t)$. Similarly, $G_t^i(X_t, U_t^i)$ is a vector of control constraints that apply to U_t^i and θ_t^i is the vector of corresponding Lagrange multipliers. As before U_t^i is the $J \times 1$ vector of crop areas for economic unit i during year t.

3.4.2 Common Pool

Now we examine the case when the groundwater system is treated as a common pool resource. Instead of maximizing the sum of profits for all of the economic units, each unit's profits are maximized separately subject to the resource constraints but not the state equation. The optimization problems are

$$\max_{U_t^i} J_t^i = \pi_t^i(X_t, U_t^i) \tag{3.42}$$

subject to (3.20). The Hamiltonians only contain the current profit term, since under a common property arrangement the farmers are assumed to ignore the future effects of their actions. For each economic unit, the current-value Hamiltonian is

$$\mathcal{H}_t^i(X_t, U_t^i) = \pi_t^i(X_t, U_t^i), \qquad \forall\, t, i \tag{3.43}$$

The augmented Hamiltonian is simply

$$\tilde{\mathcal{H}}_t^i(X_t, U_t^i, \theta_t^i) = \pi_t^i(X_t, U_t^i) + (\theta_t^i)^T G_t^i(X_t, U_t^i) \tag{3.44}$$

The first-order conditions for the common pool case are

$$\nabla_{U_t} \tilde{\mathcal{H}}_t^i(\hat{X}_t, \hat{U}_t^i, \hat{\theta}_t^i) = 0 \tag{3.45}$$

$$\hat{X}_{t+1} = A\hat{X}_t + BR(\hat{X}_t)\hat{U}_t \tag{3.46}$$

$$\hat{X}_0\text{: given} \tag{3.47}$$

$$G_t^i(\hat{X}_t, \hat{U}_t^i) \leq 0 \tag{3.48}$$

$$(\hat{\theta}_t^i)^T \geq 0, \quad (\hat{\theta}_t^i)^T G_t^i(\hat{X}_t, \hat{U}_t^i) = 0 \tag{3.49}$$

where $t = 0, \ldots, T-1$ and $i = 1, \ldots, I$. Note that the state equation is still part of the system of equations to be solved even though it is ignored by the farmers during the optimization. Also, we see that the resource constraint in Eqn. 3.48 and the corresponding Kuhn-Tucker conditions in Eqn. 3.49 are the same as in the social planner's problem.

The only difference between the social planner and common property solutions comes about through the first-order condition for $\nabla_{U_t} \tilde{\mathcal{H}}_t^i$ in Eqn. 3.45. When we substitute in the definition of the Hamiltonian, the condition becomes

$$\nabla_{U_t} \tilde{\mathcal{H}}_t^i = \nabla_{U_t^i} \pi_t^i + (\theta_t^i)^T \nabla_{U_t^i} G_t^i(X_t, U_t^i) = 0 \tag{3.50}$$

By comparing Eqn. 3.50 with Eqn. 3.41 we see that they differ by the amount $\Lambda_{t+1}^T B^i R^i(X_t)$. This is the benefit (or cost) per unit of U_t^i of the externality caused by treating the groundwater system as a common property resource. It is perhaps more intuitive to express the externality in terms of groundwater recharge. The benefit (or cost) per unit of recharge is $\Lambda_{t+1}^T B^i$.

A major policy question is how big is the cost of the externality? If the cost is significant, then some type of intervention to correct the externality may be justified. Theoretically, the externality could be eliminated by levying a tax on groundwater recharge equal to the marginal cost of the recharge, $-\Lambda_{t+1}^T B^i$. This tax would be different for each economic unit and also would change with time. Both of these characteristics would make the tax quite impractical to actually implement. An even more impractical command-and-control measure would be to directly specify the optimal crop mix or recharge rate for each economic unit at each time. More practical alternatives are a uniform tax on recharge or more uniform restrictions on land use.

3.4.3 Solution Method

The mathematical model we specified in the this chapter is a discrete-time optimal control problem. There are three main approaches for solving it: 1) numerical optimal control, 2) dynamic programming, 3) non-linear programming. We choose non-linear programming because it is capable of dealing with arbitrary constraints on the state and control variables. We use the standard optimization modeling language Gams [Brooke et al., 1997] and solve it using the non-linear programming solver Conopt2 [Drud, 1996].

In order to maintain computational feasibility, we must develop simplified representations of the hydrologic dynamics which still retain the most relevant details. The following two chapters describe our representations of the unsaturated zone and groundwater systems.

Chapter 4

State-Space Model of Groundwater Flow

4.1 Introduction

In this chapter we present a state-space model of groundwater flow in the Lower Murrumbidgee study area which can be used in the hydrologic-economic model of the previous chapter. We begin by describing our conceptual model of the hydrogeology. Next, starting from the governing partial differential equation, we derive a linear, time-invariant state-space model. We then describe our upscaling and model reduction procedure which is used to reduce the number of states in the optimization model while maintaining an accurate representation of groundwater flow. Finally, we present the spatial discretization grids and parameter values such as hydraulic conductivity and aquifer thickness.

4.2 Conceptual Model of Study Area Hydrogeology

In the Riverine Province, the sediments can be divided into three formations. Starting at surface, they are the Shepparton Formation, the Calivil Formation and the Renmark Group.

4.2.1 Shepparton Formation

The Shepparton Formation is a composite aquifer-aquitard system which is generally 40-70 meters thick, but in some areas is up to 100 meters thick. The formation is a complex system of fluvio-lacustrine sediments with significant horizontal and vertical heterogeneity. It is characterized by yellow and brown sands and variegated clays. The discontinuous nature of the sands, unlike in underlying systems, discourages significant groundwater extraction. Some useful irrigation or stock supplies are obtained from shallow spear-point systems which tap the coarser shoestring sand bodies. In many places the unit is so clayey that little horizontal connection between sand layers exists and vertical flow dominates [Evans and Kellett, 1989].

Water quality is highly variable and is a complex function of conductivity, depth to water table and the history of land and groundwater use. Recharge to the Shepparton is from infiltration of rainfall and irrigation water and percolation from surface water bodies. Most of the recharge occurs in the coarser sands of the fan-head environment where the Murrumbidgee and Lachlan Rivers enter the Murray Basin. In this recharge zone, the salinity is a few hundred mg/L TDS, rising to 20,000 mg/L in other areas. Some of the observed variation in salinity may be due to the remobilization of unsaturated-zone salts by rising water tables [Arad and Evans, 1987].

4.2.2 Calivil Formation and Renmark Group

Underlying the Shepparton Formation, the Calivil Formation sediments were deposited by the ancestral Murray River drainage system. In the highland valleys, the sediments consist of coarse sands and gravels which are incised into the underlying thin Renmark Group sediments. On the downstream plains, the sediments become progressively finer grained with a steady increase in silt and clay the farther from the highlands. The Calivil is generally confined by the low-permeability clays at the base of the Shepparton, although in some areas the basal Shepparton is as coarse as the Calivil. Near the highland valleys the hydraulic conductivity of the Calivil is very high, up to 6500 m^3/day/m [Woolley and Williams, 1978]. The proportion of clay to sand increases to the west, with an associated reduction in conductivity.

Recharge to the Calivil is from downward leakage from the overlying Shepparton Formation. Using radiocarbon dating, Drury et al. [1984] concluded that recharge to the Calivil is from the Murrumbidgee River via the Shepparton. Most of this recharge occurs where the Murrumbidgee emerges from its highland valley.

The basal Renmark Group aquifer is a confined aquifer which is virtually continuous over the study area. It consists of fluvial clay, silt, sand and minor gravel with frequent carbonaceous deposits. Within the study area, most groundwater recharge occurs in the upstream area of the alluvial fan as leakage from the Murrumbidgee River. As a result of this recharge, the groundwater salinity is lowest near the river at Narrandera and increases toward the west and away from the river.

4.3 Derivation of the Reduced-Order State-Space Model

The groundwater state equation which will be used in the optimization model is derived in the following steps: 1) we derive a continuous time state equation from the groundwater flow governing equation, 2) we discretize in space on a regular grid, 3) we define the scaling of the input and output variables, 4) we apply a model order reduction technique to reduce the number of states, and 5) we discretize in time.

4.3.1 Governing Equation

We model each stratigraphic layer in the system as a 2D leaky aquifer. If the aquifer is confined, the governing differential equation for groundwater flow is [de Marsily, 1986]

$$\frac{\partial}{\partial x}\left(T\frac{\partial h}{\partial x}\right)+\frac{\partial}{\partial y}\left(T\frac{\partial h}{\partial y}\right)=S\frac{\partial h}{\partial t}-r+q_L \tag{4.1}$$

The terms on the left-hand side of Eqn. 4.1 represent horizontal flow where $h = h(x, y, t)$ is the hydraulic head $[L]$ and T is the transmissivity $[L^2T^{-1}]$, which we have assumed to be isotropic but not homogeneous. On the right-hand side, S is the storage coefficient $[-]$, $r = r(x, y, t)$ is a source or sink $[LT^{-1}]$ (positive for aquifer recharge), and $q_L = q_L(x, y, t)$ is the leakage from the aquifer into a stream or another aquifer (positive for aquifer discharge). We must also specify appropriate initial and boundary conditions.

The source or sink flux r represents recharge from infiltration or discharge from groundwater pumping; it is exogenous to the groundwater model. The leakage, on the other hand, depends on the head in the aquifer through the relation

$$q_L = -L(h_{ref} - h) \tag{4.2}$$

where h_{ref} is the reference head in the stream or adjacent aquifer and L is the leakage conductance $[T^{-1}]$. If the leakage is to another aquifer, h_{ref} is the head in the adjacent aquifer; if the leakage is to a stream, h_{ref} is the stream stage. Once we have coupled the different aquifers, the adjacent aquifer head will be endogenous, but the stream stage is always specified as a parameter. The leakance factor can be defined as

$$L = K_z/b \tag{4.3}$$

where K_z is the vertical conductivity $[LT^{-1}]$ of the aquitard or riverbed and b is its thickness $[L]$.

If the aquifer we are modeling is unconfined, Eqn. 4.1 can still be applied but now the transmissivity depends on h, and S is set equal to the specific yield S_y. The result is the nonlinear Boussinesq equation [de Marsily, 1986]

$$\frac{\partial}{\partial x}\left(Kh\frac{\partial h}{\partial x}\right)+\frac{\partial}{\partial y}\left(Kh\frac{\partial h}{\partial y}\right)=S_y\frac{\partial h}{\partial t}-r+q_L \tag{4.4}$$

where K is the horizontal hydraulic conductivity $[LT^{-1}]$. We have assumed that the elevation of the bottom of the aquifer is zero so that h is the same as the saturated thickness. When the variations in h do not cause significant changes in the transmissivity, we can linearize this equation around a nominal transmissivity $T = Kh_0$, where h_0 is a nominal saturated thickness. In our application, this approximation is justified so we can use Eqn. 4.1 as our governing equation. In future sections, we assume that when we apply Eqn. 4.1 to an unconfined aquifer, T represents a nominal transmissivity and S has been set equal to the specific yield.

4.3.2 Spatial Discretization

We discretize the governing equation (Eqn. 4.1) in space (but not time) using the standard finite-difference approximation. Each variable and parameter is defined in space only at the N finite-difference nodes, which are arranged in a regular grid. The governing equation becomes a matrix equation

$$Mh(t) = S\frac{dh}{dt} - r(t) + L[h(t) - h_{ref}] \tag{4.5}$$

where M is the $N \times N$ finite-difference transmissivity matrix, S is a $N \times N$ diagonal matrix of storage coefficients, $h(t)$ is a $N \times 1$ vector of nodal heads, $r(t)$ is a $N \times 1$ vector of nodal recharges, L is a $N \times N$ diagonal matrix of leakage conductances, and h_{ref} is a $N \times 1$ vector of nodal heads in the adjacent aquifer or stream [de Marsily, 1986]. The elements of the finite-difference matrix M are proportional to the average transmissivities in the finite-difference cells and inversely proportional to the areas of the cells. Rearranging, we get a spatially-discrete, continuous-time state equation for groundwater head

$$\dot{h} = S^{-1}(M - L)h + S^{-1}(r + Lh_{ref}) \tag{4.6}$$

where $\dot{h}$ is a vector of head time derivatives and the time arguments have been dropped for clarity.

There will be a state equation of this form for each aquifer in the groundwater system. The state equations for the different aquifers are coupled through the leakage term. Instead of writing a system of equations, we maintain a simplified notation by assuming that the heads for all the aquifers are concatenated into a single head vector h and that the leakage conductances relating the heads in different aquifers have been incorporated into the transmissivity matrix M. This means that the only nonzero elements left in the leakage matrix L correspond to leakage to streams and not to aquifers.

By rewriting the equation in terms of deviations, we can eliminate the constant term. We define $\bar{h}$ to be the equilibrium head that results from a nominal recharge $\bar{r}$. Setting $\dot{h} = 0$ and $r = \bar{r}$, we find the equilibrium head

$$\bar{h} = -(M - L)^{-1}(\bar{r} + Lh_{ref}) \tag{4.7}$$

Using the definitions $\tilde{h} = h - \bar{h}$ and $\tilde{r} = r - \bar{r}$, the state equation (4.6) becomes

$$\begin{aligned}
\dot{\tilde{h}} &= S^{-1}(M - L)(\tilde{h} + \bar{h}) + S^{-1}(\tilde{r} + \bar{r} + Lh_{ref}) \\
&= S^{-1}(M - L)[\tilde{h} - (M - L)^{-1}(\bar{r} + Lh_{ref})] + S^{-1}\tilde{r} + S^{-1}(\bar{r} + Lh_{ref}) \\
&= S^{-1}(M - L)\tilde{h} + S^{-1}\tilde{r}
\end{aligned} \tag{4.8}$$

We are also interested in quantities besides the heads, such as the leakage to a river or other boundary. These quantities of interest are modeled in the output equation. The

leakage is

$$\begin{aligned} q_L &= L(h - h_{ref}) \\ &= L(\tilde{h} + \bar{h} - h_{ref}) \\ &= L[\tilde{h} - (M - L)^{-1}(\bar{r} + Lh_{ref}) - h_{ref}] \\ &= L\tilde{h} - L(M - L)^{-1}(\bar{r} + Lh_{ref}) - Lh_{ref} \end{aligned} \tag{4.9}$$

If we define the equilibrium leakage

$$\bar{q_L} = -L(M - L)^{-1}(\bar{r} + Lh_{ref}) - Lh_{ref} \tag{4.10}$$

then we write the equation for leakage as a deviation from equilibrium

$$\tilde{q}_L = q_L - \bar{q_L} = L\tilde{h} \tag{4.11}$$

In standard form, our linear state-space model is

$$\dot{x} = Ax + Bu \tag{4.12}$$

$$y = Cx \tag{4.13}$$

where

$$x = \tilde{h} \qquad A = S^{-1}(M - L) \tag{4.14}$$

$$u = \tilde{r} \qquad B = S^{-1} \tag{4.15}$$

$$y = \begin{bmatrix} \tilde{h} \\ \tilde{q}_L \end{bmatrix} \qquad C = \begin{bmatrix} I \\ L \end{bmatrix} \tag{4.16}$$

In systems terminology, Eqn. 4.12 is called the state equation and Eqn. 4.13 is the output equation; the variables x, u and y are called the state, input and output variables.

4.3.3 Input/Output Scaling

The state-space model derived in the previous section has state, input and output vectors of equal dimension. That is, if there are N nodes in our finite-difference grid, there will be N states (heads), N inputs (recharges) and N of each type of output (heads and leakages). However, one of the advantages of a state-space formulation is that we have the flexibility to define inputs and outputs at different scales or of different dimensions than the states. In our hydrologic-economic model this will be very useful since the inputs and outputs of economic significance are defined over larger scales than those that are appropriate for modeling hydrologic responses. We can deal with this situation within our state-space model framework by defining the states (hydraulic heads) on a finer grid than the inputs and outputs. The inputs and outputs are defined on the economic unit scale while the states are defined on the smaller hydrologic scale. The correspondence between these scales is shown in Figure 3.1.

This difference in scales means that the inputs have to be downscaled and the outputs upscaled. The scaling is done through simple linear operations defined by upscaling and downscaling matrices. To upscale the outputs, we use a weighted average of the hydrologic cells in each economic units. To downscale the inputs, we perform the inverse operation. As an illustration, suppose there are six states in the system, $x = [x_1\ x_2\ x_3\ x_4\ x_5\ x_6]^T$, and two economic units which correspond to $\{x_1, x_2\}$ and $\{x_3, x_4\}$. The states x_5 and x_6 are not part of any economic unit. There are two inputs $u^e = [u_1^e\ u_2^e]^T$ and two outputs $y^e = [y_1^e\ y_2^e]^T$, where the superscript e indicates that they are defined on the economic unit scale. We define the downscaling matrix D, so that

$$u = Du^e \tag{4.17}$$

where u is at the hydrologic cell scale. There are as many rows in D as there are hydrologic cells (n_h), and as many columns as economic units (n_e). The upscaling matrix U is used for the outputs

$$y^e = Uy \tag{4.18}$$

where y is at the hydrologic cell scale and U has dimension $n_e \times n_h$. The two operations are related by

$$U_{ij} = D_{ji}c_j, \qquad i = 1 \ldots n_h,\ j = 1 \ldots n_e \text{ (no sum over j)} \tag{4.19}$$

where c_j is the number of hydrologic cells in economic unit j.

To give a concrete example, we write out all the components and use sample values for the scaling matrices. Downscaling through Eqn. 4.17 becomes

$$\begin{bmatrix} u_1 \\ u_2 \\ u_3 \\ u_4 \\ u_5 \\ u_6 \end{bmatrix} = \begin{bmatrix} 1.2 & 0 \\ 0.8 & 0 \\ 0 & 1 \\ 0 & 1 \\ 0 & 0 \\ 0 & 0 \end{bmatrix} \begin{bmatrix} u_1^e \\ u_2^e \end{bmatrix} \tag{4.20}$$

and upscaling using Eqn. 4.18 would be

$$\begin{bmatrix} y_1^e \\ y_2^e \end{bmatrix} = \begin{bmatrix} 0.6 & 0.4 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0.5 & 0.5 & 0 & 0 \end{bmatrix} \begin{bmatrix} y_1 \\ y_2 \\ y_3 \\ y_4 \\ y_5 \\ y_6 \end{bmatrix} \tag{4.21}$$

In this case both c_1 and c_2 are 2 since there are two hydrologic cells in each of the economic units. Note that the mean of nonzero elements of each column of D in Eqn. 4.20 must equal one, and the the sum of each row of U in Eqn. 4.21 must equal one.

We redefine the system variables and matrices originally given in Eqns. 4.14 through 4.16 to give the inputs and outputs at the economic unit scale

$$x = \tilde{h} \qquad A = S^{-1}(M - L) \tag{4.22}$$

$$u = \tilde{r}^e \qquad B = S^{-1}D \tag{4.23}$$

$$y = \begin{bmatrix} \tilde{h}^e \\ \tilde{q}_L \end{bmatrix} \qquad C = \begin{bmatrix} U \\ L \end{bmatrix} \tag{4.24}$$

The actual weighting factors that define D and U are based on the amount of arable land in each hydrologic cell. The details of this are given in Section 4.4.4.

4.3.4 Model Reduction

We use a standard technique from linear systems theory to reduce the number of states in the model without significantly changing the input-output response. This is necessary so we can model the groundwater system in sufficient detail without excessive computational cost. When we are finished, the reduced number of states in the system will no longer correspond to groundwater heads, although the definitions of the inputs and outputs will not have changed.

The state-space model which we derived in the previous sections is not unique; it is only one possible realization which gives the correct input-output response. In particular, our state-space model has more states than are necessary to adequately model the system. These extra states are a computational burden. We can reduce the number of states by choosing new states which are linear combinations of the most important original states. Original states which are not significantly affected by the inputs or outputs can be eliminated. In systems theory, the amount of coupling between inputs and states is called reachability and between outputs and states is called observability. Our goal is to eliminate the states which are the least reachable and observable.

One method of doing this is known as balanced truncation, originally described by Moore [1981]. In this approach, the states of the system are transformed so that the new states are equally reachable and observable. This transformed system is called a balanced realization. Then the states which are the least reachable or observable can be truncated (deleted). The degrees of reachability and observability can be measured by the reachability Gramian

$$P = \int_0^\infty e^{\tau A} B B^T e^{\tau A^T} d\tau \tag{4.25}$$

and the observability Gramian

$$Q = \int_0^\infty e^{\tau A^T} C^T C e^{\tau A} d\tau \tag{4.26}$$

If the state transition matrix A is stable[1], then P and Q can be found by solving the algebraic

[1] A linear, time-invariant dynamical system is stable if all the eigenvalues of A are in the open left half plane; that is, $\mathrm{Re}\,\mathrm{eig}(A) < 0$.

Lyapunov equations

$$AP + PA^T + BB^T = 0 \tag{4.27}$$
$$A^TQ + QA + C^TC = 0 \tag{4.28}$$

We apply a balancing transformation T to the system matrices to give the balanced system

$$\hat{A} = T^{-1}AT \tag{4.29}$$
$$\hat{B} = T^{-1}B \tag{4.30}$$
$$\hat{C} = CT \tag{4.31}$$

The reachability and observability Grammians of the transformed system are

$$\hat{P} = T^{-1}P(T^{-1})^T = \mathrm{diag}(\Sigma_1, \Sigma_2, 0, 0) \tag{4.32}$$
$$\hat{Q} = T^TQT = \mathrm{diag}(\Sigma_1, 0, \Sigma_3, 0) \tag{4.33}$$

where $\Sigma_1, \Sigma_2, \Sigma_3$ are positive definite diagonal matrices and

$$\Sigma_1 = \mathrm{diag}(\sigma_1, \ldots, \sigma_m). \tag{4.34}$$

The zeros on the diagonal of the controllability (observability) Grammian indicate that the the corresponding states are uncontrollable (unobservable).

The σ_i are the singular values of PQ, which are defined by

$$\sigma_i = \mathrm{eig}(PQ)^{1/2} \tag{4.35}$$

These are called the Hankel singular values of the system, and they indicate the importance of the corresponding states to the input-output response. If the system has order n (i.e., there are n states), then there will be n singular values

$$\sigma_1 \geq \sigma_2 \geq \cdots \geq \sigma_m > \sigma_{m+1} = \cdots = \sigma_n = 0 \tag{4.36}$$

If $m < n$, then $n - m$ of the states in the system are unnecessary[2] and can be eliminated without changing the input-output response. We can eliminate even more states by examining the magnitudes of the singular values. The states which are least observable-reachable correspond to the smallest σ_i. To obtain a kth-order model ($k \leq m$), we partition the system matrices

$$\hat{A} = \begin{bmatrix} \hat{A}_{11} & \hat{A}_{12} \\ \hat{A}_{21} & \hat{A}_{22} \end{bmatrix}, \qquad \hat{B} = \begin{bmatrix} \hat{B}_1 \\ \hat{B}_2 \end{bmatrix}, \qquad \hat{C} = \begin{bmatrix} \hat{C}_1 & \hat{C}_2 \end{bmatrix} \tag{4.37}$$

where $\hat{A}_{11}$ is $k \times k$, $\hat{A}_{22}$ is $(n-k) \times (n-k)$ and the other matrices have appropriate dimensions. After truncation, the reduced model is defined by the system submatrices $\hat{A}_{11}$, $\hat{B}_1$ and $\hat{C}_1$.

[2]In this case, the realization is said to be nonminimal. A minimal realization would not include the uncontrollable or unobservable states and would have order m.

Moore [1981] proved that this reduced model is stable and balanced with observability and reachability Grammians are both equal to $\text{diag}(\sigma_1, \ldots, \sigma_k)$.

The theoretical error bound for the kth-order model is

$$\|e\|_2 = \|y - \hat{y}\|_2 \leq 2 \left(\sum_{i=k+1}^{n} \sigma_i \right) \|u\|_2 \tag{4.38}$$

where $\| \cdot \|_2$ is the standard Euclidean norm [Glover, 1984].

To carry out the model reduction, we used the code in the Matlab Robust Control Toolbox [Chiang and Safonov, 1998]. The numerical algorithm, described by Safonov and Chiang [1989], uses a Schur decomposition of PQ is used to avoid problems with ill-conditioning. In Section 4.6 we examine how good an approximation of groundwater response the reduced-order model provides.

4.3.5 Time Discretization

At this point our reduced-order state-space model is in continuous time. In order to embed it in the optimization model, we must discretize it in time. We do this by solving the state equation analytically. The solution to the homogeneous part of Eqn. 4.12 is

$$x(t) = e^{A(t-t_0)} x(t_0) \tag{4.39}$$

where $x(t_0)$ is the known initial state and $e^{A(t-t_0)}$ is the state transition matrix defined using a matrix exponential.[3] The particular solution is

$$x(t) = e^{A(t-t_0)} x(t_0) + \int_{t_0}^{t} e^{A(t-\tau)} B u(\tau)\, d\tau \tag{4.40}$$

If we assume $u(t)$ is constant over the interval $[t_0, t)$,[4] we can integrate to get

$$x(t) = e^{A(t-t_0)} x(t_0) + [e^{A(t-t_0)} - I] A^{-1} B u_t \tag{4.41}$$

where I is the identity matrix and u_t is the value of $u(t)$ over $[t_0, t)$.

If we discretize in time using a constant timestep Δt, we have

$$x_{t+1} = e^{A\Delta t} x_t + [e^{A\Delta t} - I] A^{-1} B u_t \tag{4.42}$$

where x_{t+1} are the states at time $t+1$, x_t are the states at time t, and u_t are the inputs which are constant over $[t, t+1)$. The output equation (Eqn. 4.13) is easily discretized as

$$y_t = C x_t. \tag{4.43}$$

[3] The matrix exponential e^{At} is defined by [Rugh, 1996]

$$e^{At} = \sum_{k=0}^{\infty} \frac{1}{k!} A^k t^k.$$

[4] In systems terminology, this is called zero-order hold.

4.4 Specification of Model Parameters

We base our specification of the model parameters on a numerical model of the entire Lower Murrumbidgee which was developed by NSW Department of Land and Water Conservation (DLWC). The original model was documented in Punthakey et al. [1994]. Recently, the DLWC updated the model by extending the calibration period and switching to a monthly timestep. The model is solved using the numerical code MODFLOW-96 [Harbaugh and McDonald, 1996].

The most recent version of the DLWC model has a gridcell size of 7.5 km $\times$ 7.5 km and three vertical layers. The model layers represent the three layers in the conceptual model: the Shepparton Formation, the Calivil Formation and the Renmark Group. The model simulates actual groundwater conditions in the Lower Murrumbidgee from 1975 to 1990 using a monthly timestep. Figure 4.1 shows the relationship between the DLWC model and our study area. The top layer is modeled as an unconfined aquifer and so gives rise to a nonlinear governing equation. The lower two layers are modeled as confined aquifers.

The DLWC model was constructed over the entire Lower Murrumbidgee watershed using a 7.5 km square grid. Our model is focused on the area surrounding the MIA with a finer grid. In addition, the DLWC model has three layers representing the Shepparton, the Calivil and the Renmark units. We have combined the Calivil and Renmark units into one model layer since they have similar hydraulic properties and small vertical gradients between them.

In order to adapt the parameters of the DLWC model to our model we must deal with two main issues: 1) interpolating model parameters to the new model domain, grid size and timestep, and 2) linearizing the governing equation parameters and boundary conditions. The follow sections describe how we derived the parameters.

4.4.1 Model Domain and Discretization

Since the study area is focused on the MIA, we can use a smaller model domain than the DLWC model. We also use different finite-difference (FD) cell sizes. For the top layer, representing the Shepparton Formation, we use a grid size of 5 km $\times$ 5 km in order to more accurately represent the depth to groundwater. This depth is critical to our model application. The heads in the lower formations are not as important, so we use a significantly coarser discretization than in the DLWC model. Instead of two vertical layers and 7.5 km $\times$ 7.5 km cell size, we use only one layer and a cell size of 10 km $\times$ 10 km.

The model grid for the top layer of the DLWC model is shown in Figure 4.1. The top layer corresponds to the top layer of our model. The middle and bottom layer of the DLWC model is shown in Figures 4.2 and 4.3. These two layers were combined into one layer in our model.

4.4.2 Hydraulic Conductivity and Storativity

We use bilinear interpolation to define the hydraulic properties at the new grid scale. Since the DLWC model assumes the top layer is unconfined, the parameters specified are the

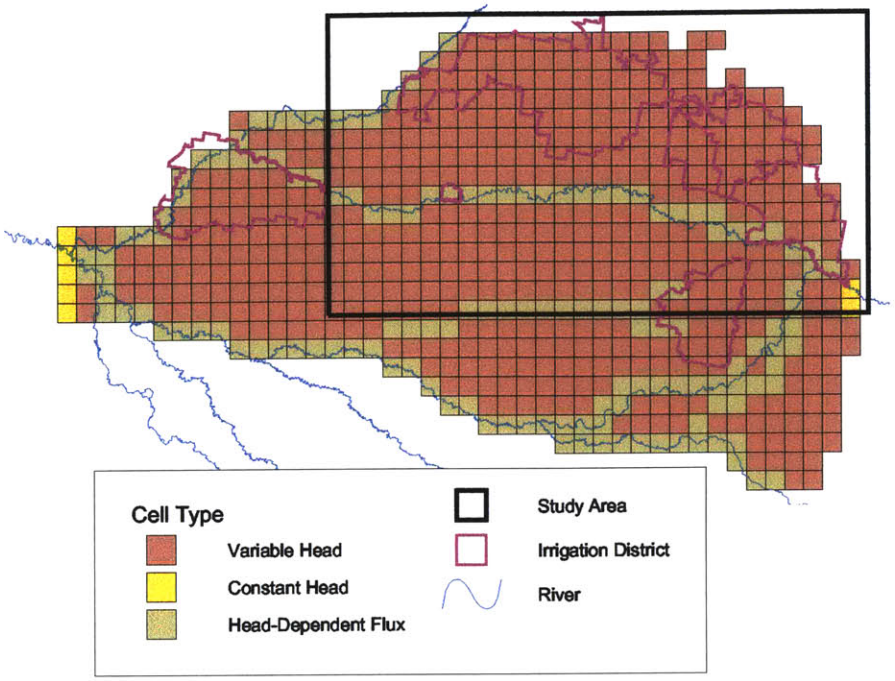

Figure 4.1: Top Layer of the DLWC Modflow Model Grid Showing Cell Type

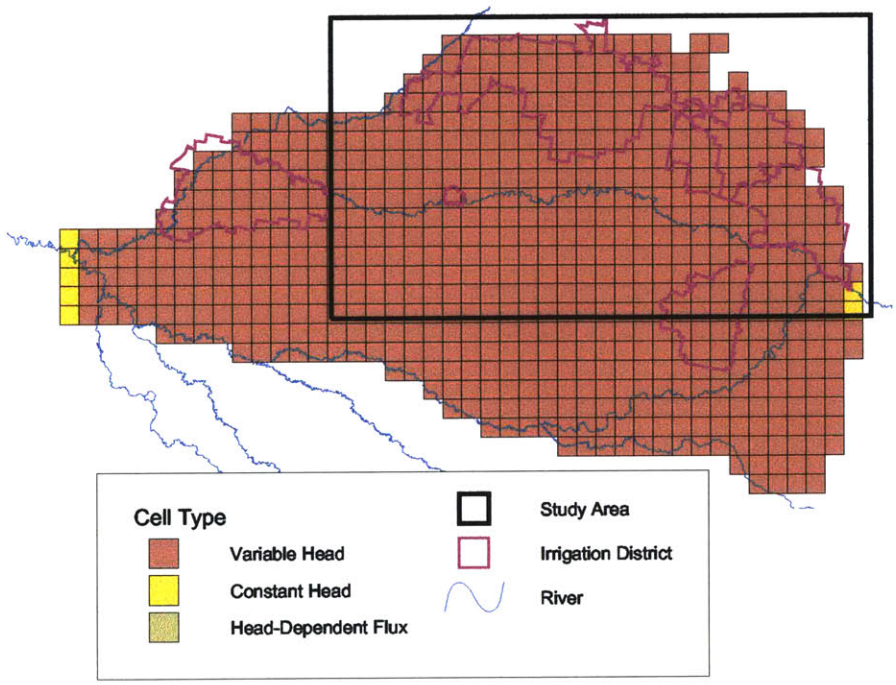

Figure 4.2: Middle Layer of the DLWC Modflow Model Grid Showing Cell Type

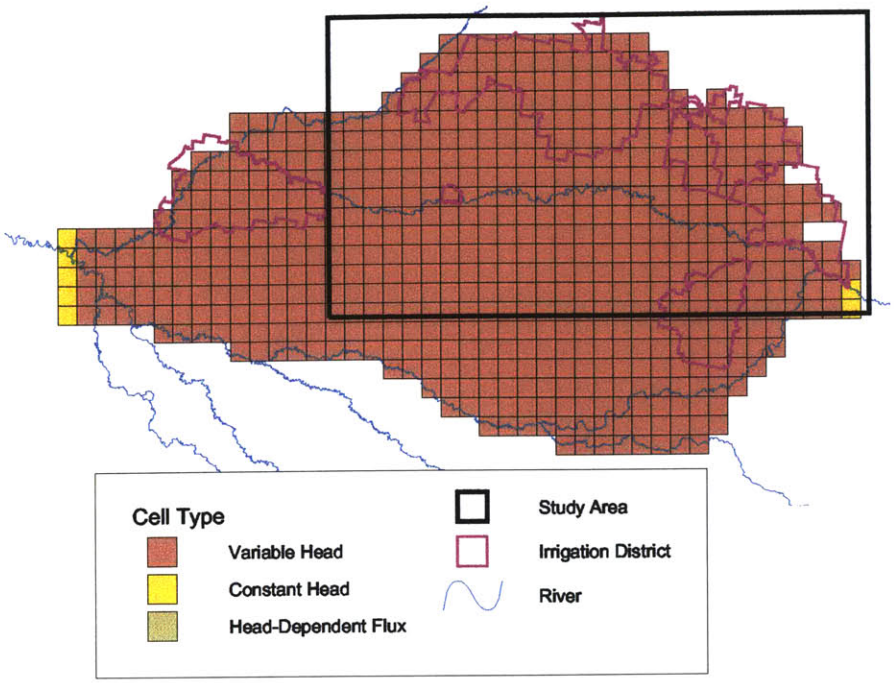

Figure 4.3: Bottom Layer of the DLWC Modflow Model Grid Showing Cell Type

hydraulic conductivity and the elevation of bottom of the Shepparton Formation. We interpolate these parameters onto the new grid. We calculate the saturated thickness of the top layer using the observed heads in 1990. By multiplying the saturated thickness by the conductivity, we get the transmissivities. To find the transmissivities for the bottom layer, we simply add the transmissivities for Calivil Formation and the Renmark Group and then interpolate onto the new grid. The resulting transmissivities for the top layer (5 km $\times$ 5 km cells) are shown in Figure 4.4 and for the bottom layer (10 km $\times$ 10 km cells) in Figure 4.5.

To estimate the vertical leakance factor between the two layers, we simple interpolate the vertical leakance from the DLWC model's top two layers. The result is shown in Figure 4.6. It has the scale of the top layer (5 km $\times$ 5 km) since it is the most detailed. We are not simulating vertical flow between the Calivil and Renmark, so the vertical leakance between these layers in the DLWC model is not used.

The specific yield from the Shepparton Formation is simply interpolated onto the new grid as shown in Figure 4.7. The storage coefficients for the Calivil and Renmark layers were averaged using the layer thicknesses as weighting factors. The result was interpolated onto the new grid. Figure 4.8 shows the final bottom layer storativities.

4.4.3 Linearization of the River Boundary Condition

In the DLWC Modflow model, river-aquifer interaction is represented using the MODFLOW-96 River Package. This package implements a head-dependent flux condition which is non-

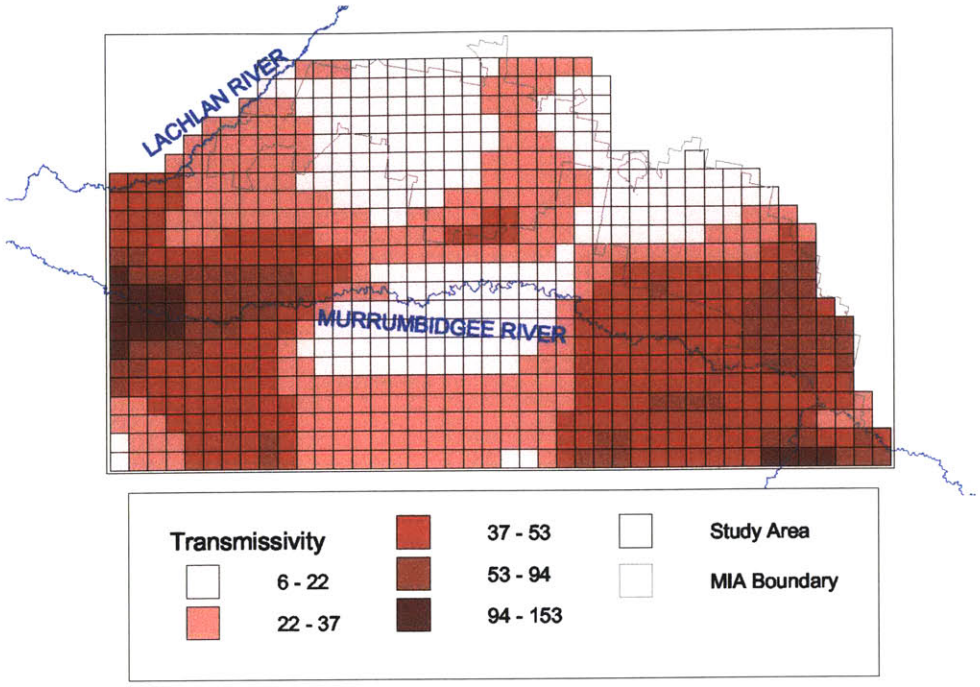

Figure 4.4: Transmissivity [m^2/day] of the Top Model Layer

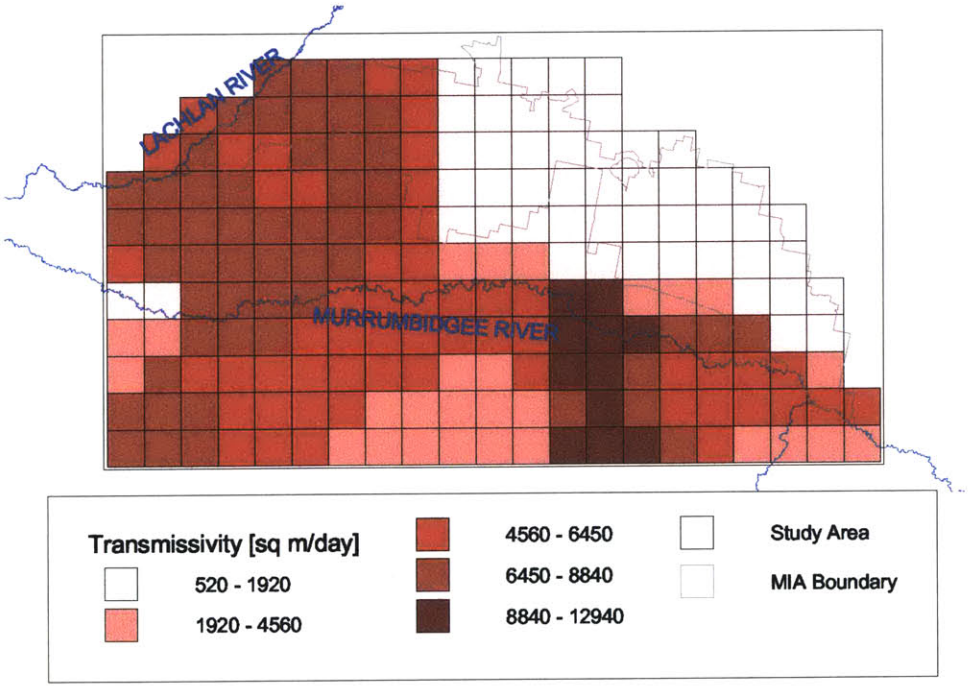

Figure 4.5: Transmissivity [m^2/day] of the Bottom Model Layer

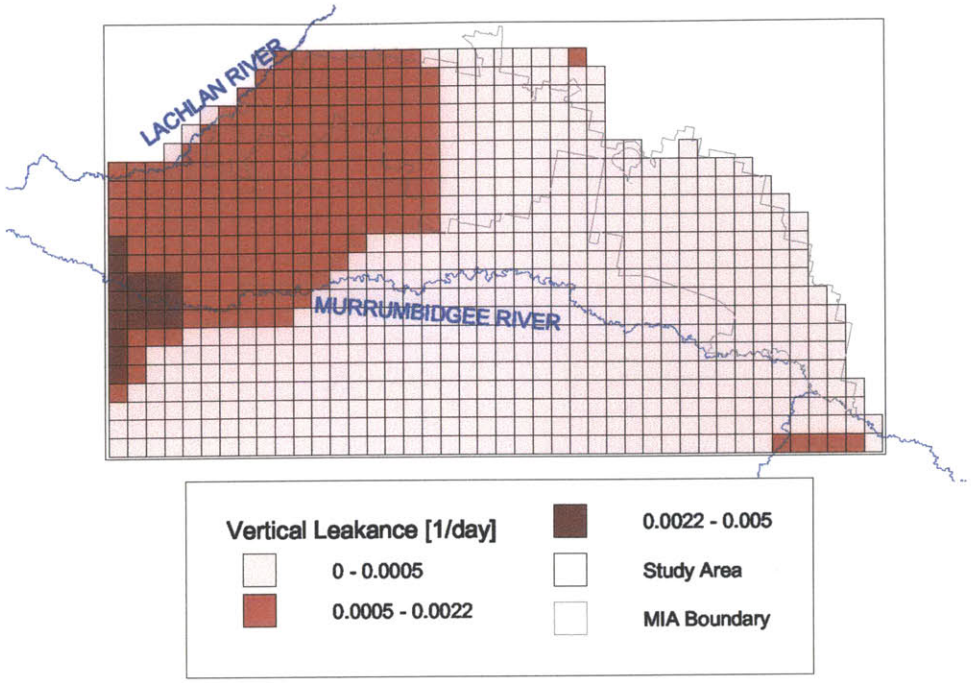

Figure 4.6: Vertical Leakance [1/day] Between the Top and Bottom Model Layers

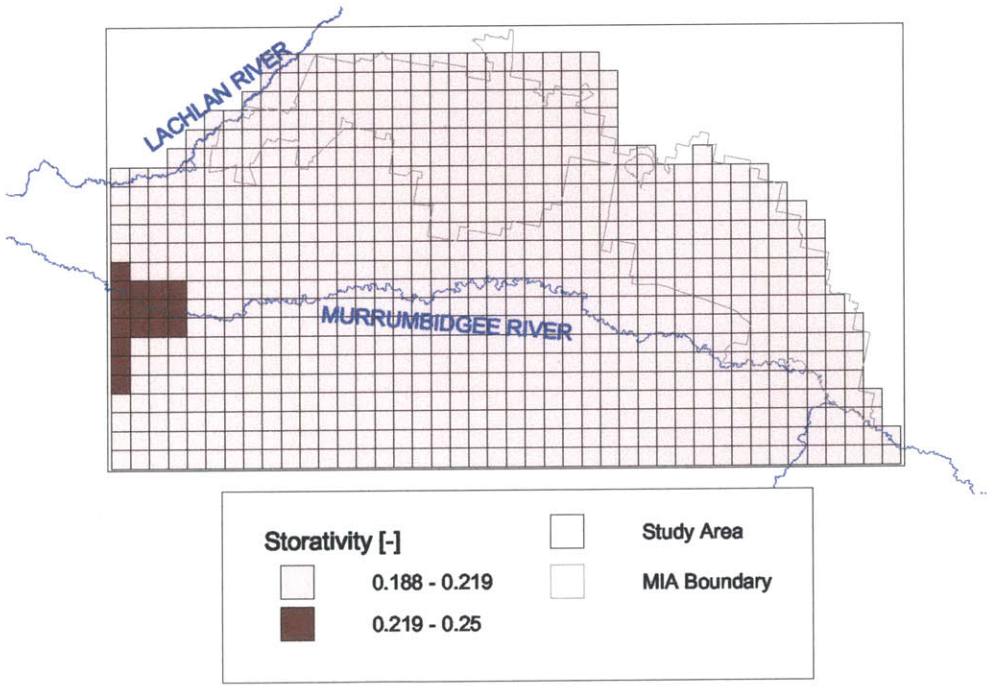

Figure 4.7: Storativity [-] of the Top Model Layer

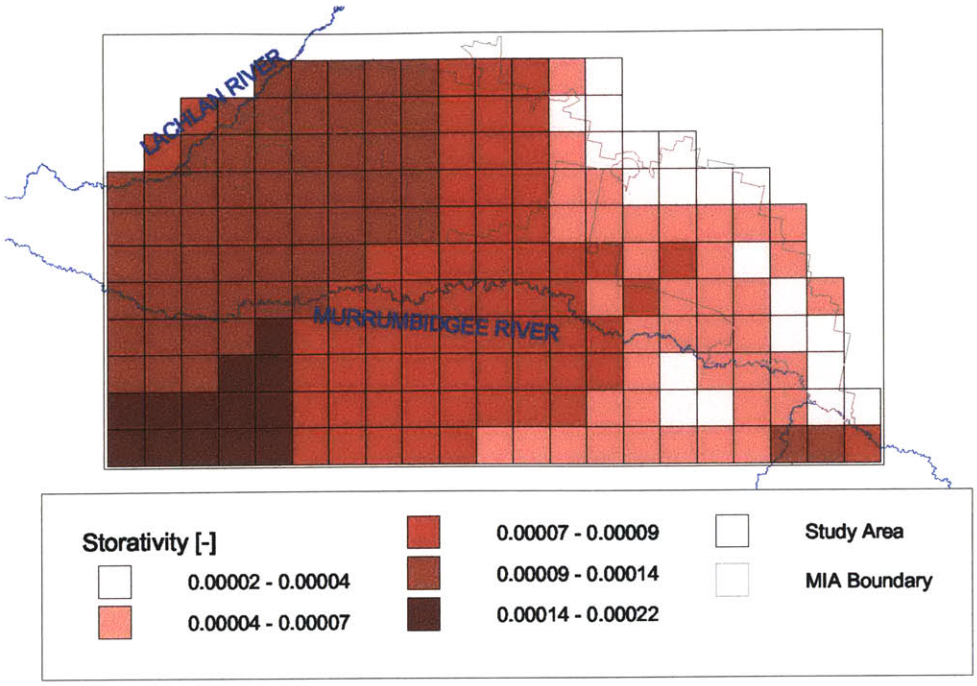

Figure 4.8: Storativity [-] of the Bottom Model Layer

linear [Harbaugh and McDonald, 1996]. The flux into the river q_{riv} is given by

$$q_{riv} = \begin{cases} L(h_{stage} - h_{aquifer}) & \text{if } h_{aquifer} > h_{bottom}, \\ L(h_{stage} - h_{bottom}) & \text{if } h_{aquifer} \leq h_{bottom}. \end{cases} \tag{4.44}$$

where L is the leakance factor, h_{stage} is the river stage, h_{bottom} is the elevation of the bottom of the river sediments and $h_{aquifer}$ is the head in the underlying aquifer.

We linearize this boundary condition by dividing the river cells into two categories: 1) specified flux and 2) linear head-dependent flux. If the aquifer head always remains below the river sediment bottom then the cell is converted to specified flux. Otherwise, the cell remains as a head dependent flux but without the cutoff. Since the heads in the model do not fluctuate a great deal, this change does not affect the model response significantly.

The area covered by the cells representing the Murrumbidgee River covers 29% less area in our model than in the DLWC model. To correct for this, we increase the leakance by 29%. For the Lachlan River, the areas representing the river in the two model are approximately the same so no correction is necessary.

4.4.4 Input and Output Scaling

In this section, we describe the procedure for finding the input/output variable scaling matrices U and D which were discussed in general in Section 4.3.3. For each economic unit, the upscaling matrix U performs a weighted average of the groundwater heads in the hydrologic cells which comprise the unit. The weighting factors for each hydrologic cell are found by

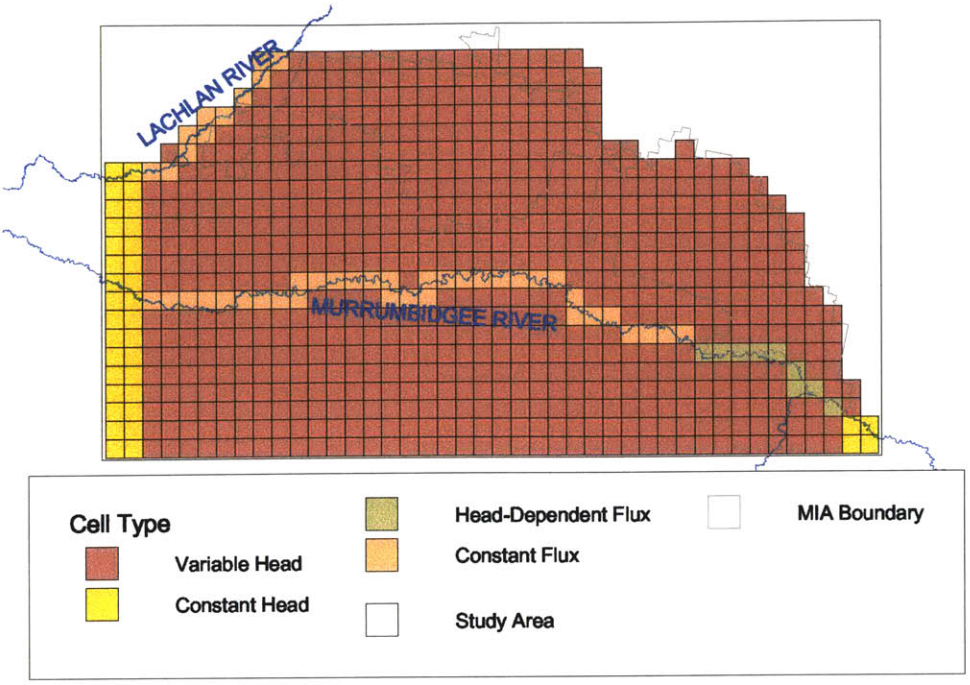

Figure 4.9: Top Model Layer Grid

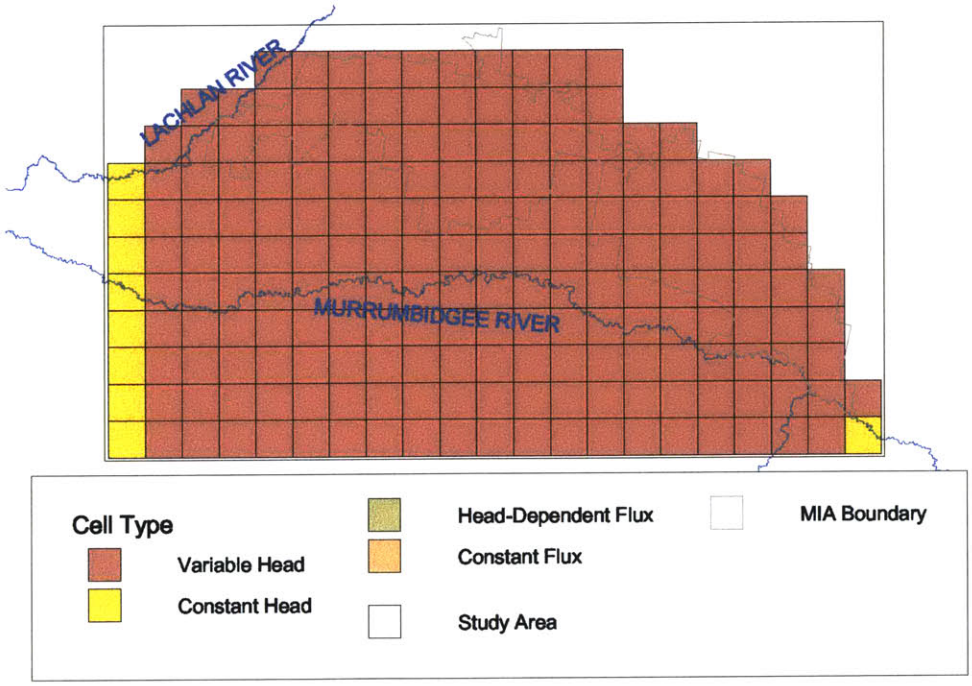

Figure 4.10: Bottom Model Layer Grid

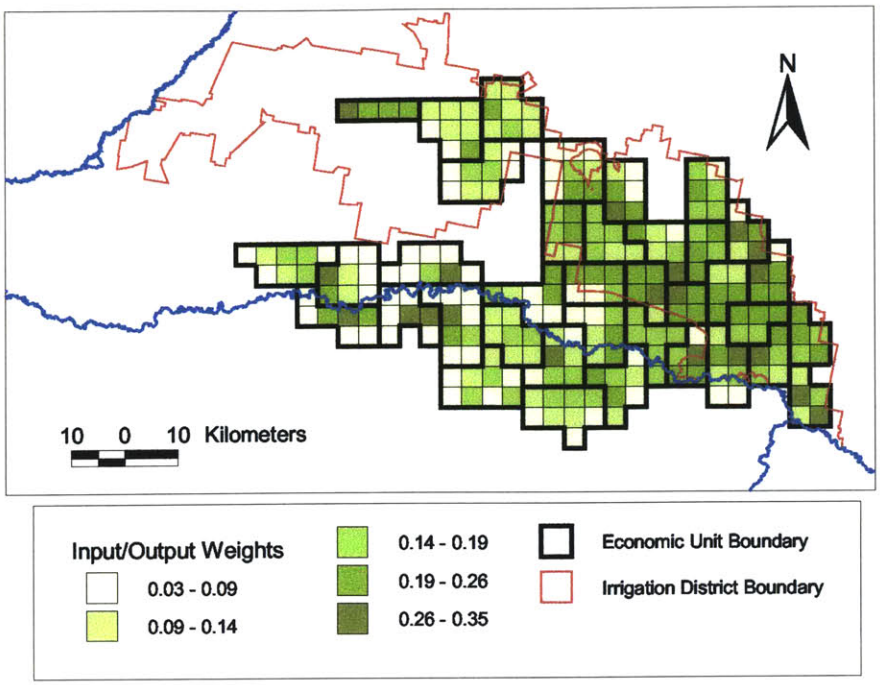

Figure 4.11: Input and Output Variables Scaling Weights

dividing the amount of arable area in each cell by the total arable area in the economic unit. We use this weighting approach since the heads in the most intensively cultivated areas have the most effect on the aggregate yields and groundwater recharge. The weighting factors for the upscaling matrix are shown in Figure 4.11. The corresponding downscaling factors are found from Eqn. 4.19.

4.5 Nominal Recharge

We assume a nominal recharge rate in areas outside of the economic units. In the top layer, we use a positive areal recharge rate of .01 mm/day for the nominal forcing outside of the economic units. Within the economic units in the top layer the nominal recharge is zero. The nominal forcing for the bottom layer is from groundwater pumping. Since the groundwater pumping decisions are exogenous to the model, we assume pumping rates are constant over time. The nominal pumping for the bottom layer is shown in Figure 4.12. These pumping rates are shown as positive values, but represent negative recharge rates.

4.6 Model Reduction Results

After reducing the model order, we were left with a model with 150 states down from an original model with 904 states. A minimal realization would have 198 states. An additional 48 states were eliminated which have little effect on the input-output response. The theoretical

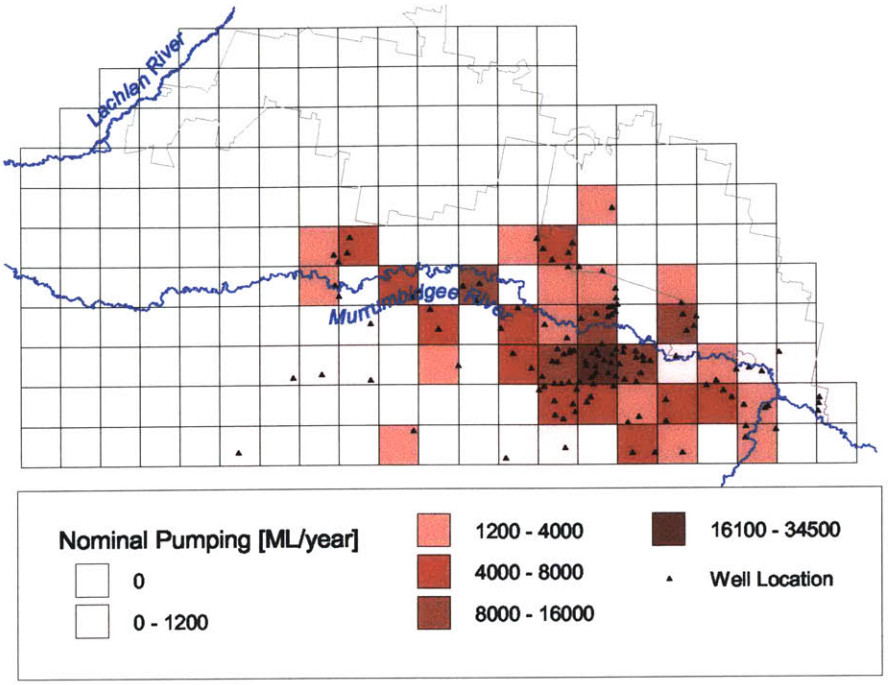

Figure 4.12: Nominal Pumping From the Bottom Layer

error bound as given by Eqn. 4.38 for our reduced order model is

$$\|y - \hat{y}\|_2 \leq 10^{-5} \tag{4.45}$$

where we have assumed that the inputs (groundwater recharges for each economic unit) are on the order of 0.1 mm/day. We can see that the sum of squared errors in the output (average heads in the economic units) is quite small. This means the reduced-order model of groundwater flow is a very good approximation of the full-order model.

Chapter 5

Model of the Unsaturated Zone System

5.1 Introduction

In this chapter, we describe the simple representation of crop production and unsaturated zone hydrology that is embedded in the optimization model described in Chapter 3. The chapter begins with a description of our general modeling approach. We then present a detailed mathematical model of the unsaturated zone system followed by a description of the model parameters applicable to the study area. This model is solved numerically for the different soil types and crops found in the study area, and the results of these detailed simulations provide the basis for estimating the simpler representations used in the optimization model. The final result is three nonlinear functions for net recharge, relative yield and irrigation water use which are specified for each crop in each economic unit. The main independent variable for these functions is depth to the piezometric surface of the Shepparton formation.

5.2 Modeling Approach

The unsaturated zone system plays a crucial role in our integrated model of salinization since crop yields and the partitioning of water fluxes are determined by complex unsaturated zone dynamics. The depth and salinity of a shallow watertable affects the soil salinity levels in the rootzone. This salinity affects the rate of crop transpiration, which is related to yield. The watertable also affects the other components of the unsaturated zone water balance: irrigation, deep percolation, and runoff. Not all of these quantities are important for the purposes of our regional-scale hydrologic-economic model. For example, we do not need to know the actual level of soil salinity. We are only interested in its effect on crop yields, groundwater recharge and irrigation water use averaged over each economic unit.

5.2.1 Simulation Experiments

We need to know these three quantities for every combination of crop, soil type, and groundwater depth and salinity. Since it is not practical to conduct field experiments for each combination, we use a detailed simulation model to provide the data for estimating simpler unsaturated zone system representations. The mathematical formulation of this model is given in Section 5.3. We conduct the simulation experiments over a thirty-year horizon using actual daily meterological forcing from 1962-1992.

5.2.2 Unsaturated Zone System Definition

There is an unsaturated zone system associated with each of the economic units. It is essentially a one-dimensional Darcy column with a top boundary at the ground surface and a bottom boundary seven meters below the surface. Below seven meters, the groundwater system begins. Through the clay soils of the unsaturated zone we are assuming vertical flow. Below this in the sandier deposits of the Shepparton Formation, the flow is assumed to be horizontal. See Figure 5.1.

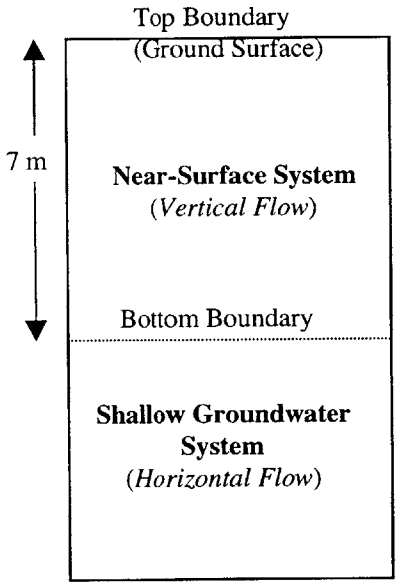

Figure 5.1: Unsaturated Zone System

5.2.3 System Inputs and Outputs

Following a systems approach, we define the inputs and outputs of our unsaturated zone system. A summary of the inputs and outputs is shown in Table 5.1. As part of the simulation experiments we vary each of the inputs in turn to determine its effect on the outputs. The inputs are kept constant throughout the 30-year simulation horizon, and the outputs are averaged over the 30-year horizon. These averaged results are then fit with simple functional forms. It is these simple representations which are used in the optimization model.

There are four unsaturated zone inputs, two of which are endogenous and change with time and two of which are exogenous and static. The endogenous dynamic inputs are the depth to the shallow groundwater piezometric surface (based on the state variables) and the crop mix (based on the control variables). The exogenous static inputs, which describe the characteristics of each economic unit, include the area of each soil type and the groundwater salinity of the shallow aquifer. We assume that groundwater salinity does not change significantly with time but it varies in space. This assumption is supported by shallow groundwater salinity measurements in the upstream MIA which showed little change from 1980 to 1998 [van der Lely, 1998].

The unsaturated zone outputs are crop yield, net groundwater recharge, and irrigation requirement. These outputs reflect the average effect of salinization on a crop over 30 years for a given set of inputs. Even though soil salinity does not appear in the system model, the salinization process is represented implicitly. We average the results over the 30 year horizon in order to represent the effect of temporal variability in meterology on salinization. This is so the system model will give results representing an average over wet and dry years. We can neglect the dynamics of salt accumulation in the unsaturated zone because it happens quickly compared with the annual timestep of our model. Figure 5.2 in Section 5.5 shows the dynamic soil salinity response. Equilibrium is reached after only a few years.

Through the optimization model, the unsaturated zone system is linked to the groundwater and crop production systems. The output for net groundwater recharge becomes the input for the groundwater system. Crop yield and irrigation requirement outputs become inputs into the economic optimization model.

Table 5.1: Unsaturated Zone System Inputs and Outputs

Name	Symbol	Note
Exogenous Inputs		
Soil Type Mix	ω_s^i	From soil map
Shallow Groundwater Salinity	c	Assumed static, spatially variable
Endogenous Inputs		
Groundwater Depth	D	From groundwater system
Crop Areas	U	Control variables
Outputs		
Relative Crop Yield	Y	Salinization effect on yield
Net Recharge Per Area	R	Input to groundwater system
Irrigation Requirement Per Area	W	Variable for rice only

5.3 Mathematical Model of Irrigated Crop Production

Models of irrigated crop production generally fall into two categories: crop-based models and soil-based models [Cardon and Letey, 1992]. Crop-based models have detailed descriptions of crop physiology, but use very simple methods of calculating water flow and solute transport. Soil-based models, on the other hand, use sophisticated descriptions of unsaturated zone flow

and transport, but typically neglect plant growth dynamics. Existing crop-based models are reviewed by Jones and Richie [1990] and soil-based models by Molz [1981].

Since an accurate representation of the hydrology of crop production is essential to modeling salinization, we use a soil-based model. The hydrology of the surface and root zones was simulated with the one dimensional water flow and solute transport code Hydrus 5.0 [Vogel et al., 1996]. Water flow is simulated using Richards' equation with a root-extraction term, and solute transport using the convection-dispersion equations. The model neglects all the factors that influence crop yield except for water stress and salinity stress. The Hydrus code solves the governing equations subject to the initial and boundary conditions using a Galerkin finite-element method using linear basis functions [Vogel et al., 1996].

5.3.1 Governing Equation for Water Flow

Water flow in variably saturated soil under isothermal conditions can be described by Richards' equation [de Marsily, 1986]. The form of the equation solved by Hydrus is

$$\frac{\partial \theta}{\partial t} = \frac{\partial}{\partial z}\left[K(h)\left(\frac{\partial h}{\partial z} - 1\right)\right] - S(h, c, z) \tag{5.1}$$

where $\theta(h)$ is the volumetric water content $[L^3L^{-3}]$, h is the pressure head $[L]$, S is the root water uptake rate $[T^{-1}]$, c is the salinity concentration $[ML^{-3}]$, z is the spatial coordinate which is positive downward $[L]$, t is time $[T]$, and $K(h)$ is the unsaturated hydraulic conductivity function $[LT^{-1}]$.

Soil Hydraulic Properties

The unsaturated soil hydraulic properties which must be specified are $K(h)$, the hydraulic conductivity function, and $\theta(h)$, the soil water retention function. We assume that hysteresis can be neglected. Different functional forms have been used including those of Brooks and Corey [1966] and van Genuchten [1980]. We use a modified version of the van Genuchten relationships that have more flexibility in the description of hydraulic properties near saturation [Vogel and Cislerova, 1988]. This flexibility improves numerical performance when we simulate heavy soils. The soil water retention and hydraulic conductivity functions are:

$$\theta(h) = \begin{cases} \theta_r + \dfrac{\theta_m - \theta_r}{(1 + |\alpha h|^n)^m} & h < 0 \\ \theta_s & h \geq 0 \end{cases} \tag{5.2}$$

and

$$K(h) = \begin{cases} K_s K_r(h) & h < 0 \\ K_s & h \geq 0 \end{cases} \tag{5.3}$$

where

$$K_r = \left(\frac{\theta - \theta_r}{\theta_s - \theta_r}\right)^{0.5}\left[\frac{1 - F(\theta)}{1 - F(\theta_s)}\right]^2 \tag{5.4}$$

$$F(\theta) = \left[1 - \left(\frac{\theta - \theta_r}{\theta_m - \theta_r}\right)^{1/m}\right]^m \tag{5.5}$$

$$m = 1 - 1/n, \qquad n > 1 \tag{5.6}$$

These hydraulic functions require the specification of the five standard van Genuchten parameters (θ_r, θ_s, α, n, and K_s) and an additional parameter (θ_m) which is set slightly larger than θ_s.

Crop Root Water Uptake

The root water uptake rate depends on the potential transpiration rate and the amount of water and salinity stress the plant is experiencing. The potential transpiration rate T_p is established by the climate and crop being grown. It is an input to the model. The root-extraction term was defined by Feddes et al. [1978] as

$$S(h, c, z) = a(h, c)S_p(z), \qquad 0 \leq a(h, c) \leq 1 \tag{5.7}$$

where $a(h, c)$ is the water and salinity stress function [-] and $S_p(z)$ is the potential root water uptake rate $[T^{-1}]$. The distribution of the potential uptake over the root zone is given by $b(z)$, so that

$$S_p(z) = b(z)\, T_p \tag{5.8}$$

where $\int b(z)$ over the soil column is one. Combining these definitions gives the actual root water uptake,

$$S(h, c, z) = a(h, c)\, b(z) T_p. \tag{5.9}$$

Water and salinity stress are assumed to be multiplicative effects so that

$$a(h, c) = a_w(h) a_s(c) \tag{5.10}$$

where $a_w(h)$ is the water stress function and $a_s(c)$ is the salinity stress function. Water stress is given by

$$a_w(h) = \frac{1}{1 + (h/h_{50})^{p_1}} \tag{5.11}$$

where h_{50} is the pressure head at which the extraction rate is reduced to 50% of the potential rate and p_1 is an empirical parameter. There was no data available on appropriate values for these parameters. We assumed that $p_1 = 3$ and $h_{50} = -5000$ so that water stress would have an insignificant effect on the results.

Similarly, salinity stress is given by

$$a_s(h) = \frac{1}{1 + (c/c_{50})^{p_2}} \tag{5.12}$$

where c_{50} and p_2 are empirical parameters. When Equation 5.12 was fit to crop salt tolerance data for a large number of crops, p_2 was found to be approximately 3 [van Genuchten and Hoffman, 1984; van Genuchten and Gupta, 1993]. The values for c_{50} for each crop are described in Section 5.4.3.

5.3.2 Initial and Boundary Conditions for Water Flow

In order to solve Eqn. 5.1 the initial distribution of pressure head is specified:

$$h(z,t) = h_i(z), \qquad t = t_0 \tag{5.13}$$

The effect of an arbitrary initial condition is minimized by running the simulation for 10 years in order to calculate a realistic initial condition.

In addition to the standard specified pressure head (Dirichlet type) and specified flux (Neumann type) boundary conditions, we also used a system-dependent atmospheric boundary condition at the soil-atmosphere interface. This atmospheric boundary condition was used at the top boundary for all simulations except during the rice growing season. Rice is grown under flooded conditions, so we used a specified pressure head boundary condition set at the ponding depth of 15 cm.

The potential flux of the atmospheric boundary is controlled only by atmospheric conditions; however, the actual flux depends also on the prevailing soil moisture conditions. During a numerical simulation, the boundary condition may change between specified-head and specified-flux depending on the state of the system. This atmospheric-type boundary condition satisfies the following conditions [Neuman et al., 1974]:

$$|q| \leq |E_p| \tag{5.14}$$

$$h_A \leq h(z,t) \leq 0, \qquad z = 0 \tag{5.15}$$

where E_p is the maximum potential rate of infiltration or evaporation under the current atmospheric conditions, h is the pressure head at the soil surface, and h_A is the minimum pressure head allowed. While the pressure head is within the limits of 5.15, the flux is specified and is equal to the potential flux. When one of limits of 5.15 is reached, a specified-head condition is applied with the head equal to the limit. The minimum pressure head h_A can be determined from the equilibrium conditions between soil water and atmospheric water vapor [Feddes et al., 1974]. In our simulations, h_A was fixed at -10,000 cm. The maximum pressure head was set at 0. This is equivalent to assuming that water cannot pond on the soil surface; instead, there is instantaneous runoff.

At the bottom boundary we always apply a specified head condition. The pressure head specified is equal to the computational domain length (7 meters) minus the depth to the Shepparton (upper layer) groundwater piezometric surface.

5.3.3 Governing Equations for Salt Transport

Solute transport is described by the standard convection-dispersion equation (CDE). We assume that the salt is an inert, nonadsorbing solute. The CDE for this case is

$$\frac{\partial \theta c}{\partial t} = \frac{\partial}{\partial z}\left(\theta D \frac{\partial c}{\partial z}\right) - \frac{\partial qc}{\partial z} - SC_S \tag{5.16}$$

where c is the solute concentration in solution $[ML^{-3}]$, D is the dispersion coefficient $[L^2T^{-1}]$, q is the Darcy fluid flux density $[LT^{-1}]$, θ and S are the same as in Equation 5.1, and C_S is

the concentration of soil solution extracted by the roots $[ML^{-3}]$. The dispersion coefficient D is defined as

$$D = \epsilon|q|/\theta + D^0\tau \tag{5.17}$$

where ϵ is the dispersivity $[L]$ of the soil, D^0 is the ionic or molecular diffusion coefficient in free water $[L^2T^{-1}]$, and τ is a tortuosity factor. We use the tortuosity relationship of Millington and Quirk [1961]:

$$\tau = \frac{\theta^{7/3}}{\theta_s^2} \tag{5.18}$$

where θ is the volumetric water content and θ_s is the volumetric water content at saturation.

5.3.4 Initial and Boundary Conditions for Salt Transport

Equation 5.16 is solved subject to the initial condition

$$c(z, t_0) = c_i(z) \tag{5.19}$$

The boundary condition at the soil surface is a third-type condition:

$$\left(-\theta D\frac{\partial c}{\partial z} + qc\right)\Big|_{z=0} = \begin{cases} q_0 c_0(t) & q_0 > 0 \\ 0 & q_0 \leq 0 \end{cases} \tag{5.20}$$

where c_0 is the salt concentration of the infiltrating water and q_0 is the Darcian flux density at the soil surface. Note that q is taken as positive into the soil domain, and so at the surface it is positive downward. During periods of upward water flux, the salt flux is zero. At the bottom boundary, we modified the Hydrus code to implement the following boundary condition:

$$\begin{cases} \frac{\partial c}{\partial z}\big|_{z=l} = 0 & q(l,t) > 0 \\ c(l,t) = c_l(t) & q(l,t) \leq 0 \end{cases} \tag{5.21}$$

where c_l is the salt concentration of groundwater flowing into the soil domain. Equation 5.21 describes a zero-gradient boundary condition during downward flow and a specified-concentration condition during upward flow.

5.4 Simulation Parameters for the Study Area

In this section we specify the values of the model parameters which we used to describe the study area.

5.4.1 Soil Types and Hydraulic Properties

In the study area there are four type of soils which are used for cropping: 1) transitional red-brown earth, 2) self-mulching clay, 3) non-self-mulching clay, and 4) prior stream deposits. The distribution of these soil types is shown on Figure 2.4. As described above, we will simulate each soil type separately and then create the unsaturated zone transfer functions for each economic unit as linear combinations of the functions for the soil types. In this section, we describe the hydraulic properties used in the simulations of each of the soil types.

There have been several studies measuring the hydraulic conductivity of transitional red-brown earths (TRBE). Beecher [1991] described a TRBE at the CSIRO's field station in Whitton, in the middle of the MIA. The soil at this site is Mundiwa clay loam, an alkaline soil with 10 cm of clay loam topsoil overlying a heavy clay B horizon. It is typical of TRBE soils used for growing rice. The infiltration under flooded rice was measured during late March at 5.6 mm/day. The regional water table under the site was between 3.5 and 4 m below the surface. Total water consumption over the rice season was about 18 ML/ha. This suggests a B horizon saturated conductivity of about 6.5 mm/day. Sides et al. [1993] measured the saturated hydraulic conductivity of another TRBE soil (Willbriggie clay loam) under a rice field in Willbriggie, 12 km south of Griffith. The saturated hydraulic conductivity of a 1.5 m deep soil profile ranged from 4.8 mm/day to 21.7 mm/day, increasing with depth. The recharge rate was estimated as 0.7 mm/d with a depth to watertable of 0.8 m. Meyer et al. [1990] estimated the saturated conductivity of Mundiwa clay loam as 100 mm/day in the A horizon (0–30 cm) and 2 mm/day in the B horizon.

Other studies have only reported infiltration rates under different crops. van der Lelij [1990] estimated average percolation under rice in the region as 400 mm/season in areas with deep watertables and less than 100 mm/season in areas with shallow watertables. Self-mulching clay soil had an infiltration rates of 3.5 mm/day and TRBE had an infiltration rate of 0.9 mm/day. With a watertable at a depth of 0.85 m, upward movement during the wheat growing season was 1 to 3 mm/day with self-mulching clay and 0.3 to 1.5 mm/day for TRBE.

Our approach has been to adjust the saturated hydraulic conductivity values in order to match reported measurements and infiltration rates. We divide the soil profile into A, B, and C horizons and assign the conductivities given in Table 5.2. In addition, the bottom 2 meters of the column represents the top 2 meters of the Shepparton Formation, the shallow groundwater system. We assume this portion of the Shepparton Formation has a saturated conductivity of 0.1 m/day. For the other hydraulic properties, we use only two sets of properties, one corresponding to a sandy-loam soil and the other to a clay soil. The values are shown in Table 5.3.

5.4.2 Potential Evapotranspiration and Precipitation

Potential evapotranspiration (ET) rates are required to provide boundary conditions for the numerical simulations. This is the rate of ET which would occur with no limit on water availability. We follow the approach of Doorenbos and Pruitt [1977] and estimate ET_c, the

Table 5.2: Saturated Hydraulic Conductivity of Soil Horizons

	Saturated Hydraulic Conductivity [mm/day]		
Soil Type	A horizon (0–30 cm)	B horizon (30–200 cm)	C horizon (200–500 cm)
Non-Self-Mulching Clay	20	2	2
Transitional Red-Brown Earth	20	3	3
Self-Mulching Clay	50	5	5
Prior Stream Deposits	20	10	100

Simulated soil column for all soil types also includes 2 meters of the Shepparton formation below the C horizon with a conductivity of 100 mm/day.

Table 5.3: Modified van Genuchten Soil Parameters[a]

Soil Texture	θ_r	θ_s	θ_m	α	n
Sandy-Loam[b]	0.068	0.4100	0.4100	0.075	1.89
Clay[c]	0.068	0.3800	0.3803	0.008	1.09

[a] See Table 5.2 for hydraulic conductivity values.
[b] Used for A horizon of all soil types, C horizon of prior stream deposits, and the Shepparton formation.
[c] Used for B and C horizon of all soil types, except for C horizon of prior stream deposits.

potential ET for a particular crop or non-cropping land use, as

$$ET_c = K_c ET_o \tag{5.22}$$

where K_c is the crop/land-use factor and ET_o is reference ET. During a simulated crop growing season, the potential ET is used to specify the potential transpiration rate T_p, and the measured precipitation is used to specify the potential top flux E_p. Outside of the growing season, the potential transpiration rate is zero, and the potential ET is subtracted from precipitation to specify the potential top flux.

The reference ET for the case study area is calculated from pan evaporation measurements taken by the Australian Bureau of Meteorology at a weather station in Griffith for the period 1962-1992. Precipitation measurements from this period were also used. Monthly crop/land-use factors were taken from Wu et al. [1998] and are given in Table A.1.

5.4.3 Crop Salinity Tolerance and Applied Water Salinity

Irrigation water and soil water salinity is commonly measured using electrical conductivity meters. These measurements are reported in deci-Siemens per meter (dS/m) or micro-Siemens/cm (μS/cm). The later is also referred to as electrical conductivity or EC units. For water of moderate salinity and typical chemical composition, EC units can be converted

Table 5.4: Soil Salinity Yield Loss Parameters

Crop	Salinity Threshold[a] [dS/m]	Yield Loss[a] [-]	C_{50}[b] [mg/L][c]
Rice	4.0	0.12	5227
Wheat	2.9	0.13	4318
Annual Pasture	1.3	0.15	2965

[a] Source: NSW Department of Land and Water Conservation [1998a].

[b] Salinity which reduces yield by 50%, calculated from previous columns for use in Hydrus simulations.

[c] 1 mg/L = 640 dS/m

to mg/L total dissolved solids (TDS) by multiplying by 0.64. At salinities greater than or equal to 1500 EC (or 960 mg/L TDS), options for consumptive use of the water become restricted. The upper limit for irrigation water under normal conditions is 2300 EC (1472 mg/L TDS). The recommended limit for drinking water is 800 EC.

The water supplied to farms in the MIA for irrigation is of very good quality, around 100 to 200 EC. For the simulation we assume it is 150 EC or 96 mg/L. Rainfall is assumed to have a negligible salinity.

The salt tolerance of crops is typically reported using the linear-threshold model of Maas and Hoffman [1977]. Crop salinity tolerance data adjusted to the conditions of the MIA were collected and are summarized in Table 5.4. Hydrus requires as an input c_{50}, the concentration which reduces yield by 50%. We convert the electrical conductivity units into mg/L TDS and calculate the 50% yield concentration using the linear-threshold model. The resulting value of c_{50} is given as the final column in Table 5.4.

5.4.4 Other Parameters

Other parameters used in the simulations are shown in Table 5.5.

Table 5.5: Hydrus Simulation Parameters

Parameter	Symbol	Value
Computational domain depth	-	500 cm
Finite element size	-	1 cm
Timestep	-	0.001–0.1 day
Molecular diffusion coefficient	D^0	0.5 cm^2/day
Dispersivity	ϵ	4 cm
Max. concentration of root extraction	C_S	0.002 mg/L

5.5 Detailed Dynamic Modeling Results

In this section, we show some of the results of the Hydrus simulations in order to illustrate the unsaturated zone dynamics. These dynamic effects are averaged out in the static system model of the unsaturated zone.

In Figure 5.2, we see a simulated time series for rootzone soil salinity under 20 years of rice cultivation for different groundwater depths. We can clearly see the salinization process occuring. When the groundwater depth is less than 1 meter, the rootzone salinity levels begin to rise from their initial values. After around 15 years, the levels reach an approximate equilibrium. For the deeper groundwater depths, no salinization occurs because the water and salt flux from the bottom of profile is sufficient to prevent any accumulation of salt in the rootzone.

Figures 5.3 and 5.4 show typical profiles of soil salinity and pressure head while growing rice with a shallow groundwater head of 50 cm below the groundsurface. The first figure shows profiles during the rice growing season when the top boundary has a specified head of 10 cm to simulate ponded conditions. The second figure shows profiles during the fallow season when the top boundary is an atmosperic type. In both cases, the bottom boundary has a specified pressure head of 50 cm, and the groundwater salinity is 3,840 mg/L.

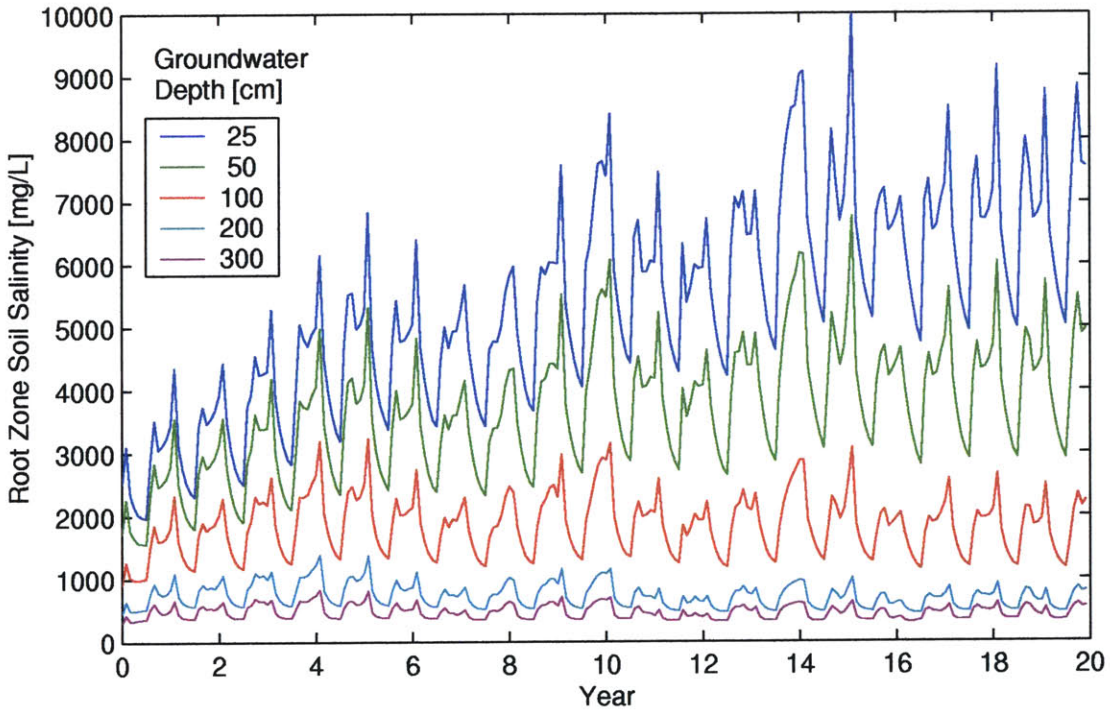

Figure 5.2: Dynamic Soil Salinity Response Under Rice With Different Watertable Depths

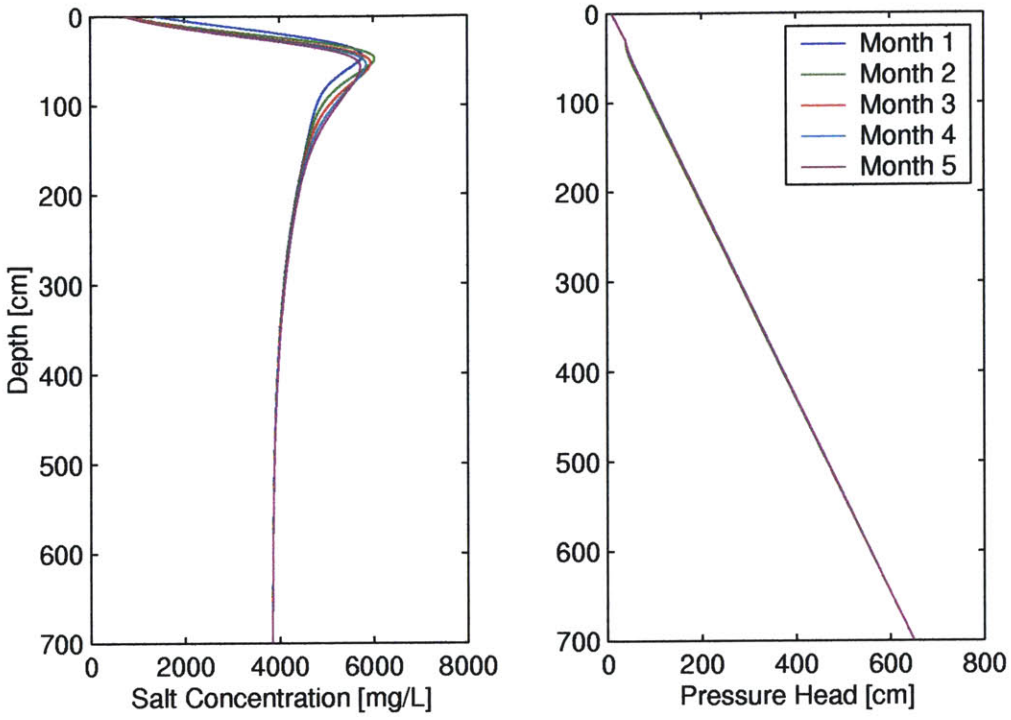

Figure 5.3: Salt Concentration and Pressure Head Profiles During Rice Growing Season

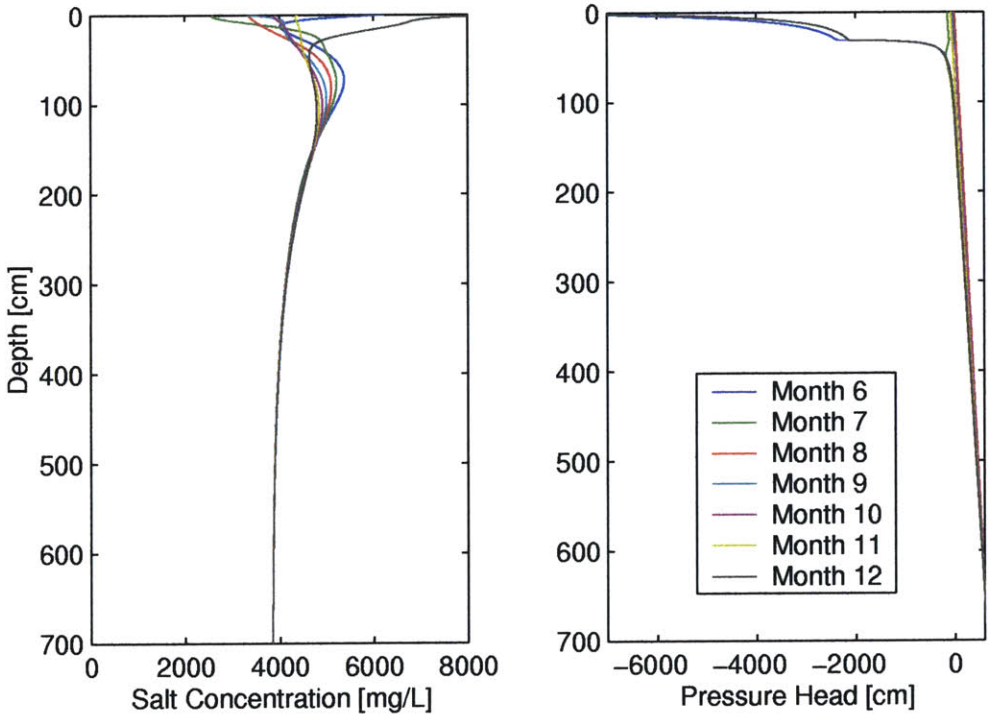

Figure 5.4: Salt Concentration and Pressure Head Profiles During Rice Fallow Season

5.6 Unsaturated Zone Transfer Function Model

Our model of the unsaturated zone system comprises the three equations relating the system inputs to the three outputs. The same functional form is used for all three equations, but the parameters vary depending on the output, the economic unit and the crop.

5.6.1 Estimation of Transfer Function Parameters For Each Soil Type

We run Hydrus simulations of each crop on each soil type at different groundwater depths and salinities to generate data describing the response of the unsaturated zone system outputs to the unsaturated zone system inputs. Each simulation was run for 40 years using actual meterological forcing. The first 20 years of data were not used in order to allow the system to reach an equilibrium soil salinity. The final 20 years were averaged so that each simulation provided one data point for each system output.

The transfer functions were fit using nonlinear least squares to the following equation

$$z_s = (\beta_1 + \beta_4 c^{\beta_5}) \exp(\beta_2 D) + \beta_3 \tag{5.23}$$

where z_s is the dependent variable (recharge, relative yield, or irrigation) for soil type s, β_i are the fitted parameters, c is the groundwater salinity concentration and d is the groundwater depth. The production function for recharge from fallow has β_4 and β_5 set equal to zero since the groundwater concentration has no effect. The simulation data and fitted curves are shown in Appendix A.2.

5.6.2 Upscaling the Transfer Functions to the Economic Units

To find the transfer functions for each economic unit, we calculate a weighted average of the functions for each soil type using the percentages of the soil types in the units as the weighting factor. Mathematically, this can be written

$$Z_i = \sum_s \omega_s^i z_s, \qquad \sum_s \omega_s^i = 1, \ \forall i \tag{5.24}$$

where Z_i is the transfer function for economic unit i, and ω_s^i is the fraction of the area of economic unit i which has soil type s.

Up to this point we have assumed that the yield per hectare is constant over each economic unit. In reality, yields may decline as more land is planted to a given crop due to land heterogeneity [Howitt, 1995]. The best land will be used first, so that the average yield will decrease as more of the lower quality land comes into production. At the same time, the costs of production may increase as more land is planted due to restricted management or machinery capacity. The result is that the net revenue per hectare or gross margins decline as more land is put into production.

We model this effect by using a factor which reduces yields as more of each crop is grown. The factor is based on a study of rice productivity across the MIA done by Harrison and

Chapman [1999]. We assume the results can be applied in the same way to each economic unit and to crops other than rice. The yield of each crop is multiplied by

$$\gamma(u_t^{ij}) = 1 - 0.6(u_t^{ij}/U_{max}^i)^3 \tag{5.25}$$

where u_t^{ij} is the land planted to each crop and U_{max}^i is the maximum land available in the economic unit i. The expression for γ is used in the definition of large-scale yield in Eqn. 3.7.

5.6.3 Groundwater Salinity

The groundwater salinity input which enters into the unsaturated zone transfer functions is from the formation immediately below the column defining the unsaturated zone system. There are extensive measurements from shallow piezometers of salinities within the MIA, but not many measurements outside it. Due to this lack of data, we assign the groundwater salinities to the economic units by zone. The best quality groundwater is in the Outside MIA zone due to the flushing effect of recharge from the river. The next best quality is in the Upstream MIA zone and the worst in the Downstream MIA zone. The values assigned are given in Table 5.6. See Figure 3.1 for the zone locations.

Figures 5.5 and 5.6 show shallow groundwater salinities in EC units for 1980 and 1998. Although there is some redistribution of salinity concentrations between the two years, there does not appear to be either an upward or downward trend over time. The spatially averaged salinity concentration was 6200 μS/cm in 1980 and 5950 μS/cm in 1998. This data supports our assumption that groundwater salinities can be treated as constant over the planning horizon of the case study.

Table 5.6: Shallow Groundwater Salinity by Zone

Economic Unit Zone	Salinity [mg/L]
Outside MIA	1,280[a]
Upstream MIA	3,840[b]
Downstream MIA	6,400[b]

[a] Source: Lawson and Webb [1998]

[b] Source: van der Lely [1998]

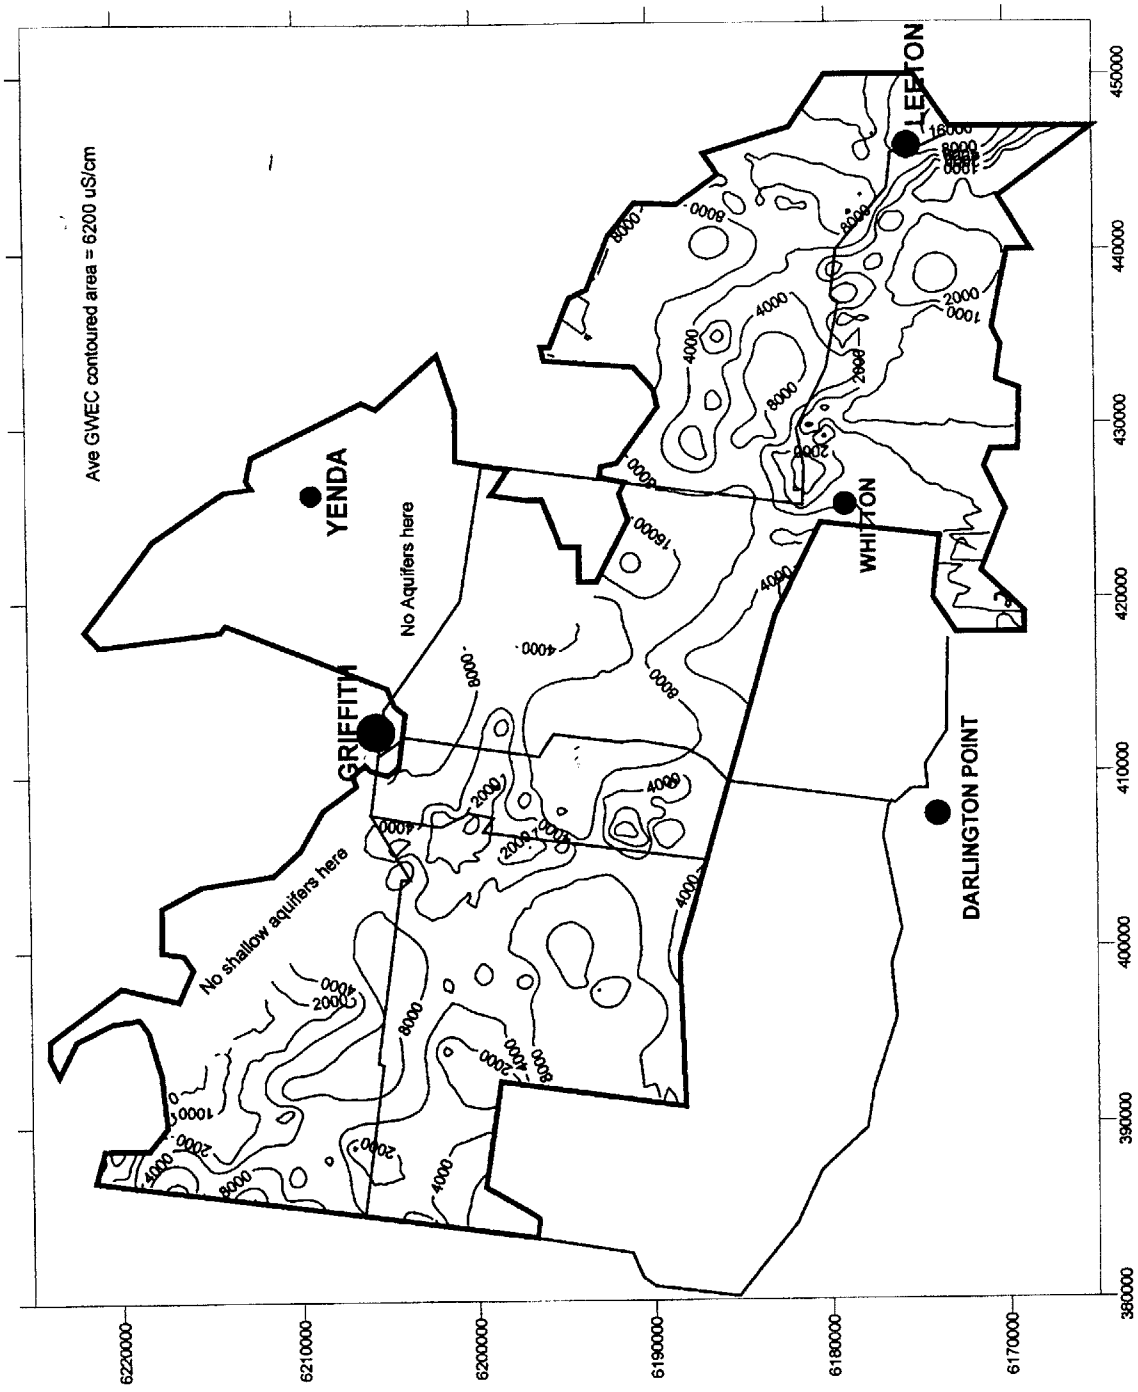

Figure 5.5: Upstream MIA Shallow Groundwater Salinity 1980
Adapted from van der Lely [1998]

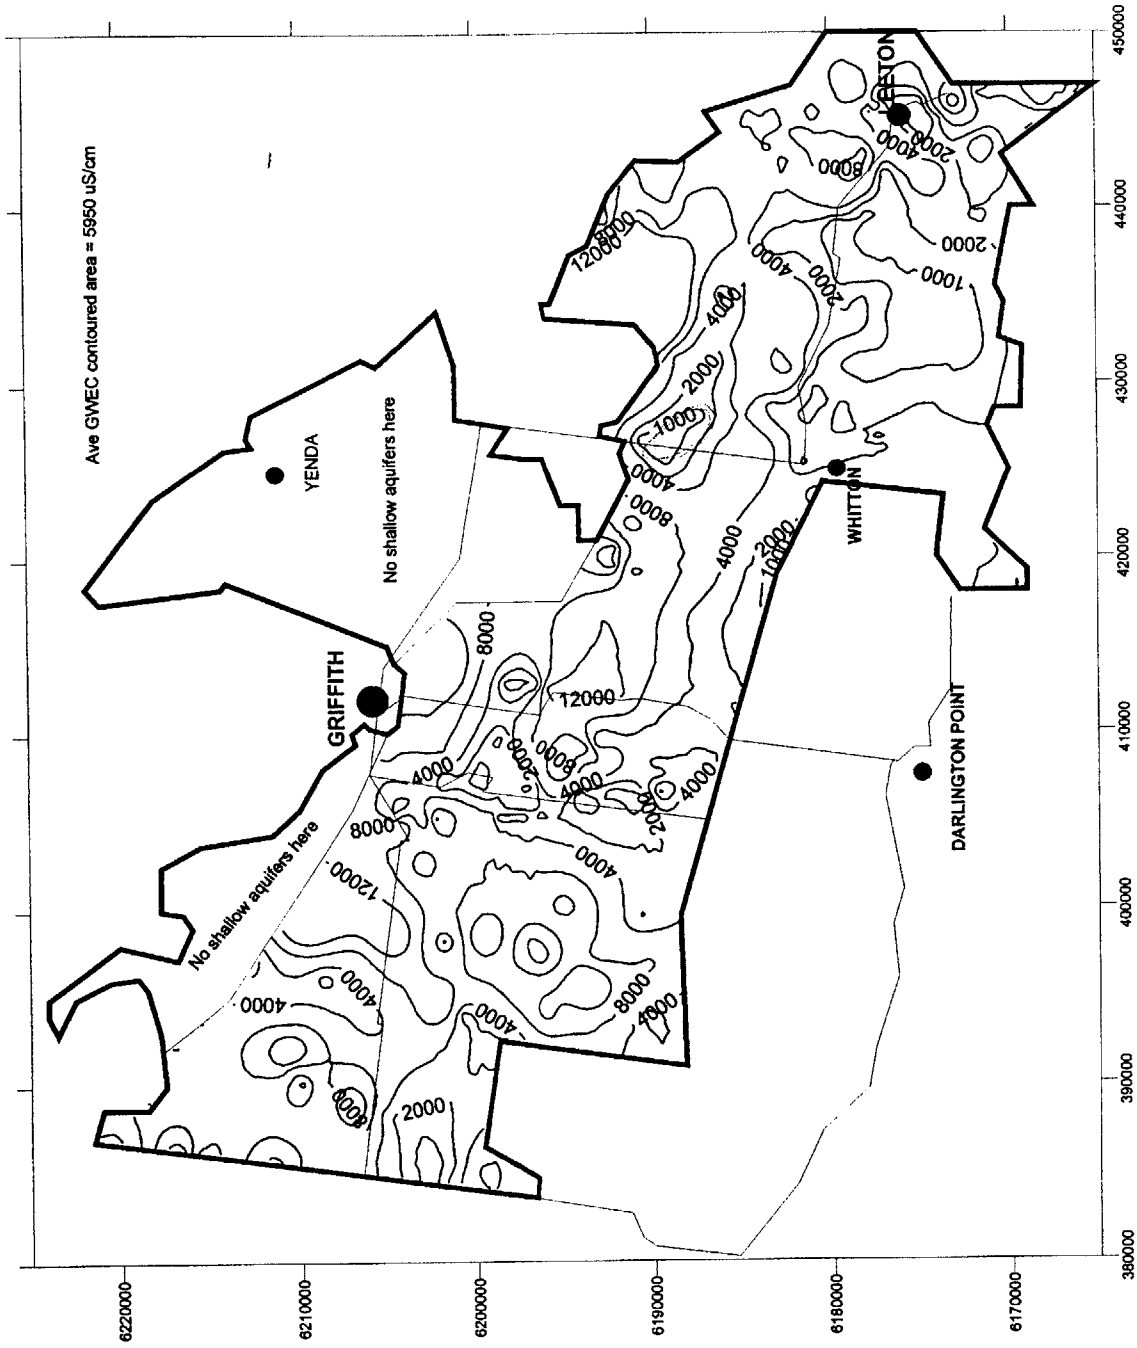

Figure 5.6: Upstream MIA Shallow Groundwater Salinity 1998
Adapted from van der Lely [1998]

Chapter 6

Results

6.1 Introduction

In this final chapter, we present the results of the scenario runs of our hydrologic-economic model of the Lower Murrumbidgee. We begin with a description of the scenarios. We then examine the cost of the common-pool externality and the benefits of the two policy options: the rice restriction and the water market. Finally, we examine the long-term behavior of the system and implications for sustainability.

6.2 Scenario Summary

The 10 scenario runs are summarized in Table 6.1. They can be divided into three groups: 15-year runs comparing the social optimum with the common pool (Scenarios 1–4), 30-year runs examining policy options (Scenarios 5–8), and 150-year runs exploring the long-term system behavior (Scenarios 9–10).

6.2.1 Scenarios 1–4: Social Optimum vs. Common Pool

In order to quantify the cost of the externality caused by the common pool groundwater system, we compare runs for the social optimum and common pool scenarios defined earlier. As explained in Chapter 3, the social optimum scenario represents the case in which cropping decisions are made in a socially optimal way such that the regional returns to agriculture are maximized. The common pool scenario represents the case in which each economic unit acts both myopically and selfishly, maximizing current revenues without regard for the future or other economic units. Due to the computational cost of simulating the social optimum solution, we use a time horizon of 15 years.

Table 6.1: Hydrologic-Economic Model Scenario Runs

Scenario Number	Social Optimum (SO)/ Common Property (CP)	30% Rice Restriction?	Water Market?	Length [years]
1	SO	no	yes	15
2	CP	no	yes	15
3	SO	no	no	15
4	CP	no	no	15
5	CP	yes	yes	30
6	CP	yes	no	30
7	CP	no	yes	30
8	CP	no	no	30
9	CP	yes	no	150
10	CP	no	yes	150

6.2.2 Scenarios 5–8: Policy Options Under Common Pool

We are interested in evaluating two policy options: the 30% rice area restriction and the surface water market. All four combinations of these policy options are run. We do this under the common-pool arrangement because this is the situation that requires policy intervention due to potential market failures. Since the model of the common pool is computationally much less costly, we use a planning horizon of 30 years.

The case with the rice restriction and a water market is the closest scenario to the present situation. The scenario with the rice restriction and no water market represents the situation before 1996 when trading of water began. A potential future arrangement is the case with the water market and no rice restriction.

6.2.3 Scenarios 9–10: Long-term Behavior

In the final two scenario runs, we examine how the system approaches steady-state. Low groundwater transmissivities lead to a groundwater system with a large time constant. It takes well over 100 years for the system to reach a steady state. We use the common pool regime and compare the most constrained case (with the rice restriction and no water market) to the least constrained case (without the rice restriction and with a water market).

6.3 Cost of the Common-Pool Externality

In the first set of scenario runs, we estimate the cost of the externality caused by the common pool property arrangement for the groundwater system. We do this by comparing the total benefits of the social optimum scenario with the total benefits of the common pool scenario. The difference is the cost of the common-pool externality. If the cost is significant, then this would provide a justification for some sort of government intervention, such as the rice

restriction. To isolate the effect of the externality, we do not include the rice restriction in these scenarios. We examine the externality cost in two cases: with and without a water market.

6.3.1 Results With Water Market

We begin with the case which includes a quasi-market for surface water. The results show little difference between the net revenues of the social optimum scenario and the common pool scenario. Figure 6.1 shows the total net revenues over time. This is a very conservative analysis because it only allows changes in crop mix. If other economic adjustments were included in the model, we would expect even less of a difference between the two cases.

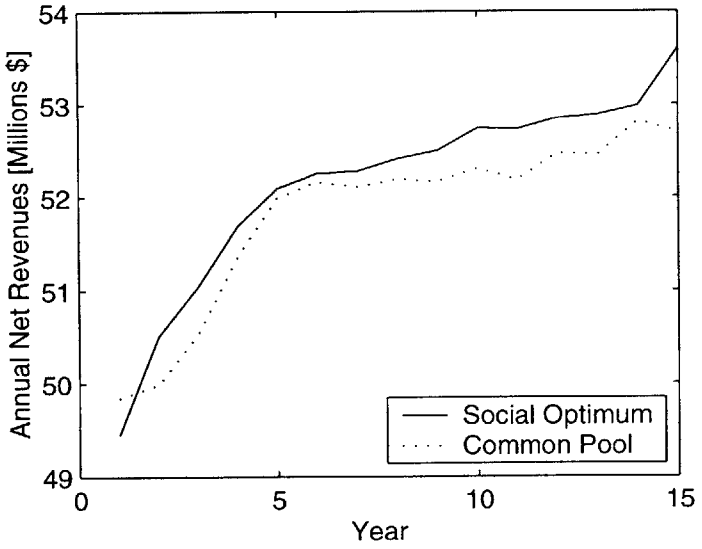

Figure 6.1: Total Net Revenues With Water Market (Scenarios 1 and 2)

The total levels of crop production are shown in Figure 6.2 and crop areas in Figure 6.3. The spatial distribution of crop areas in the economic units for the social optimum case is shown in Figures 6.4, 6.5 and 6.6. The distributions for the common pool case are not shown since they are very similar.

The average depth to groundwater under the social optimum and common pool regimes is shown in Figure 6.7. This figure also shows the difference in depths between the social optimum and common pool scenarios. A plot of the spatial distribution of watertable depths at selected times for the social optimum case in shown in Figure 6.8.

The spatial distribution of the shadow price of groundwater recharge is shown for four times in Figure 6.9. This is the shadow price for the social optimum; the shadow price for the common pool case will always be zero since the future value of the resource is ignored. The marginal cost or benefit of the common-pool externality is given by the difference between the shadow prices in the social optimum case and those in the common-pool case. Since the common-pool shadow prices are zero, Figure 6.9 also shows this difference.

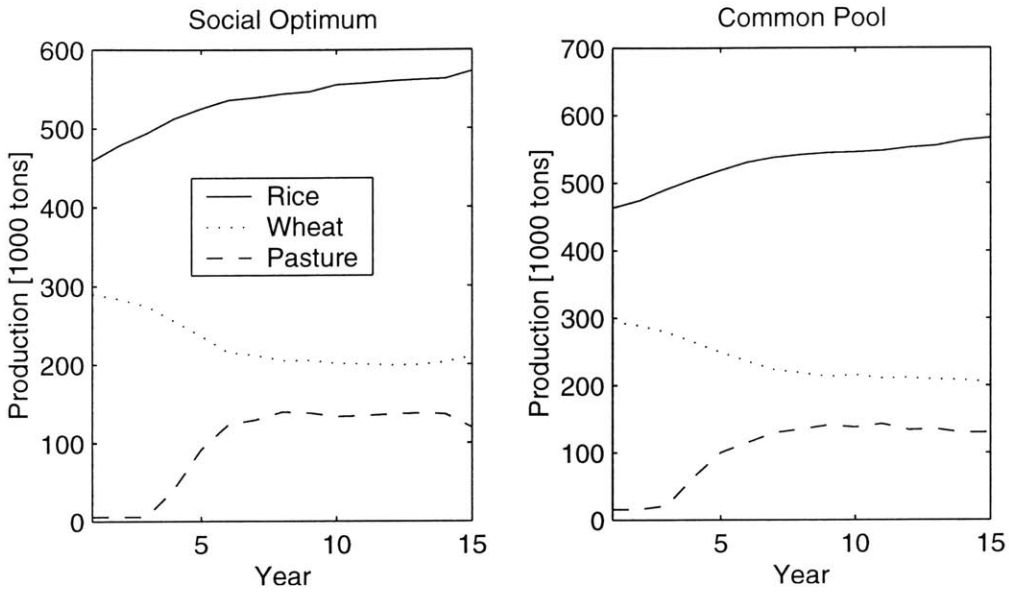

Figure 6.2: Total Crop Production With Water Market (Scenarios 1 and 2)

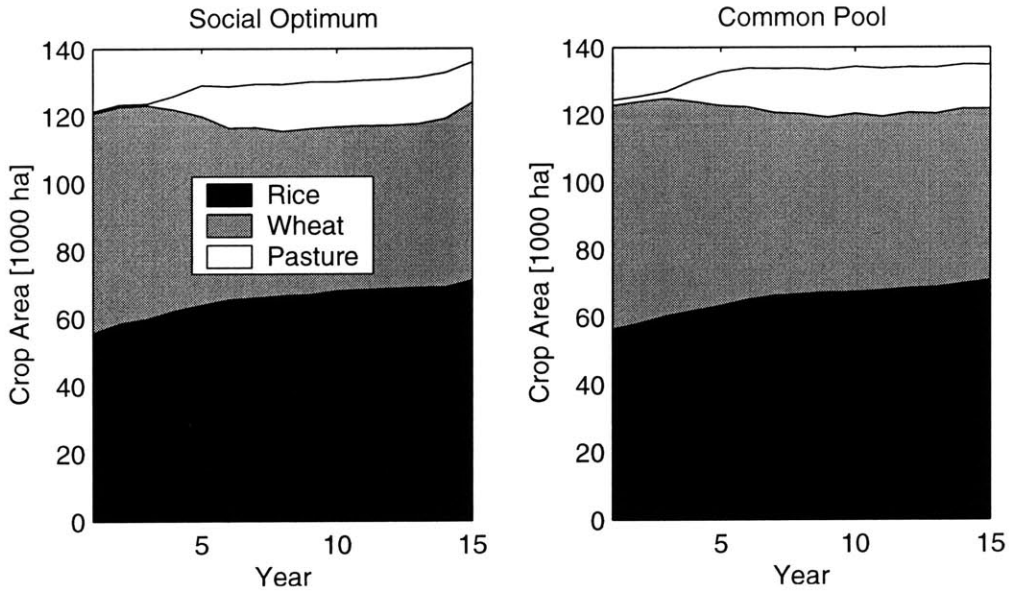

Figure 6.3: Total Crop Areas With Water Market (Scenarios 1 and 2)

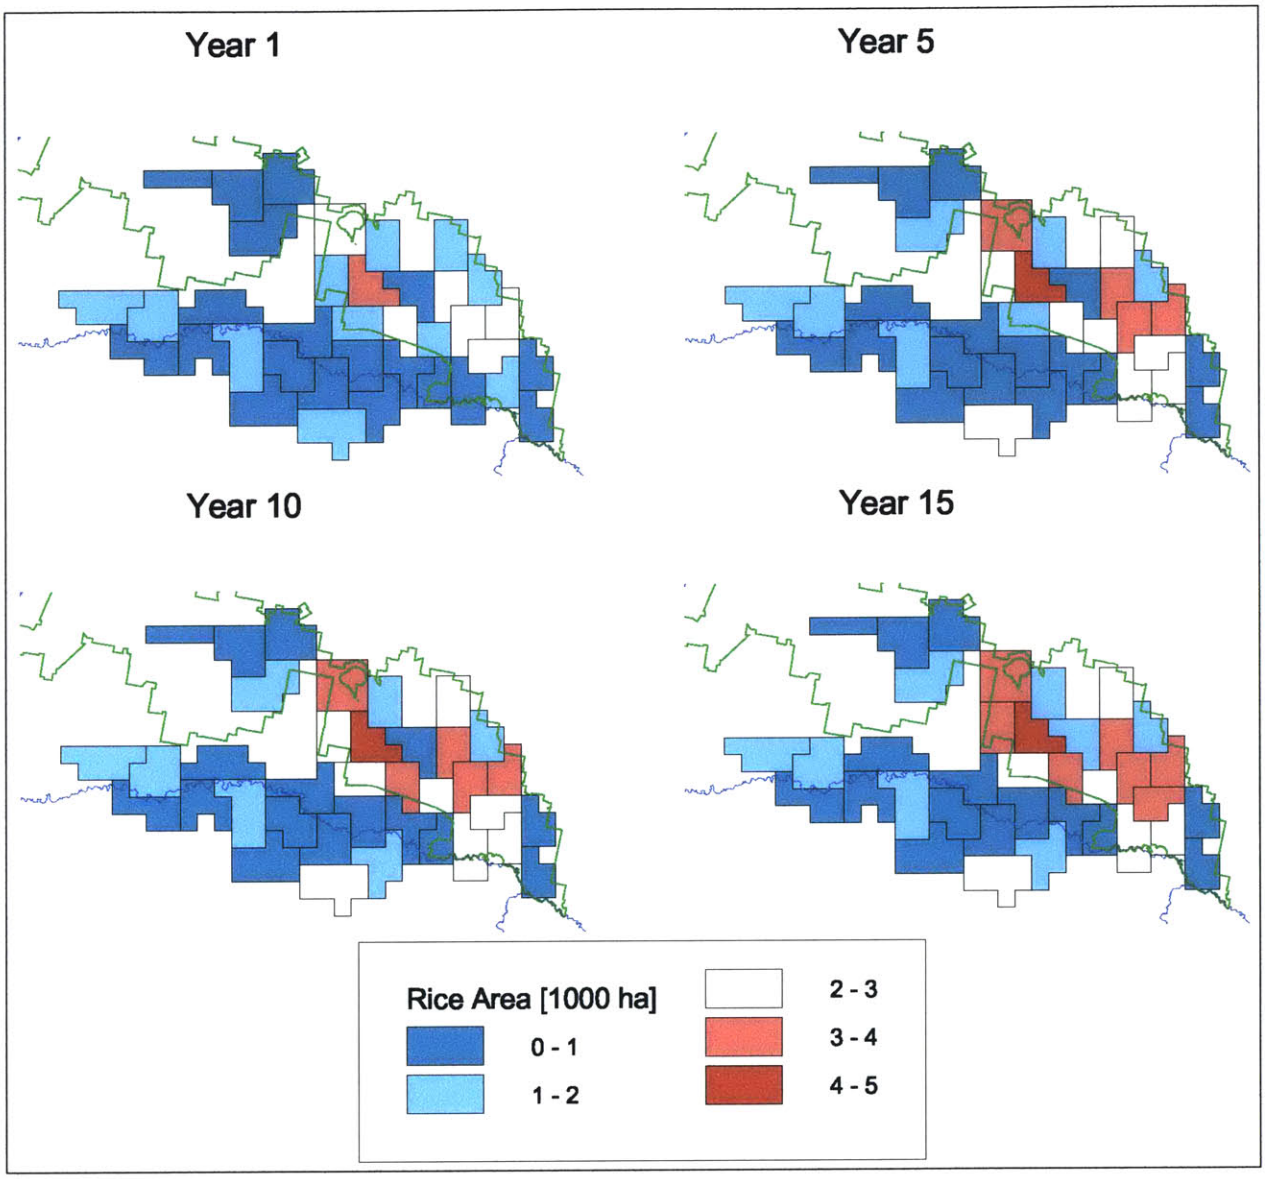

Figure 6.4: Rice Area Distribution at Selected Times (Scenario 1)

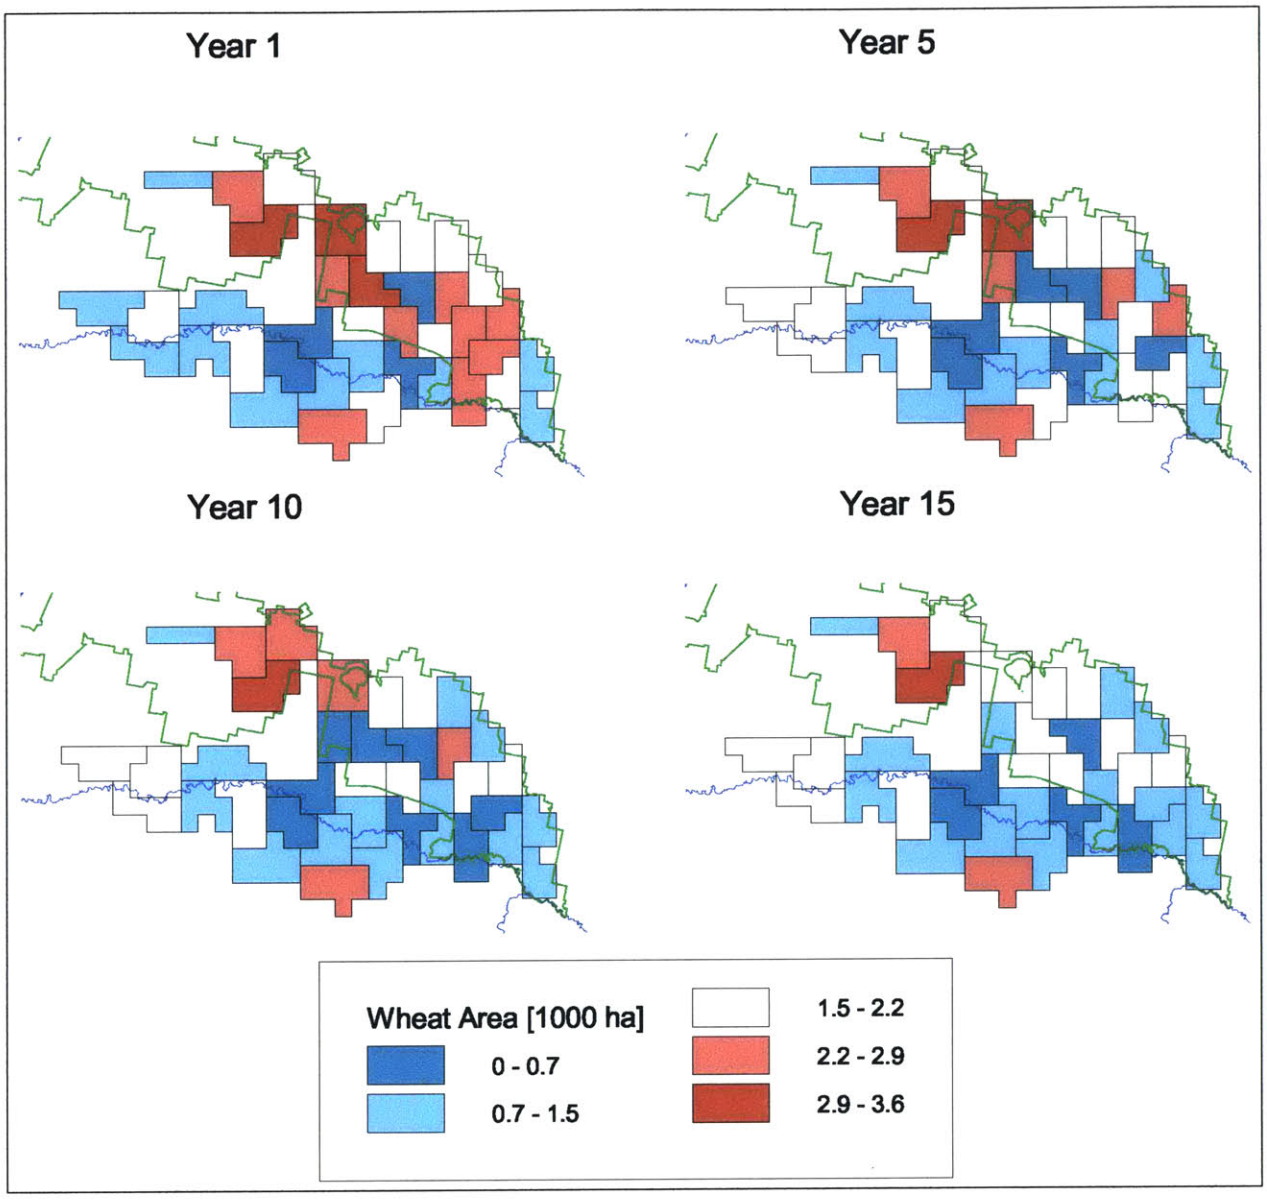

Figure 6.5: Wheat Area Distribution at Selected Times (Scenario 1)

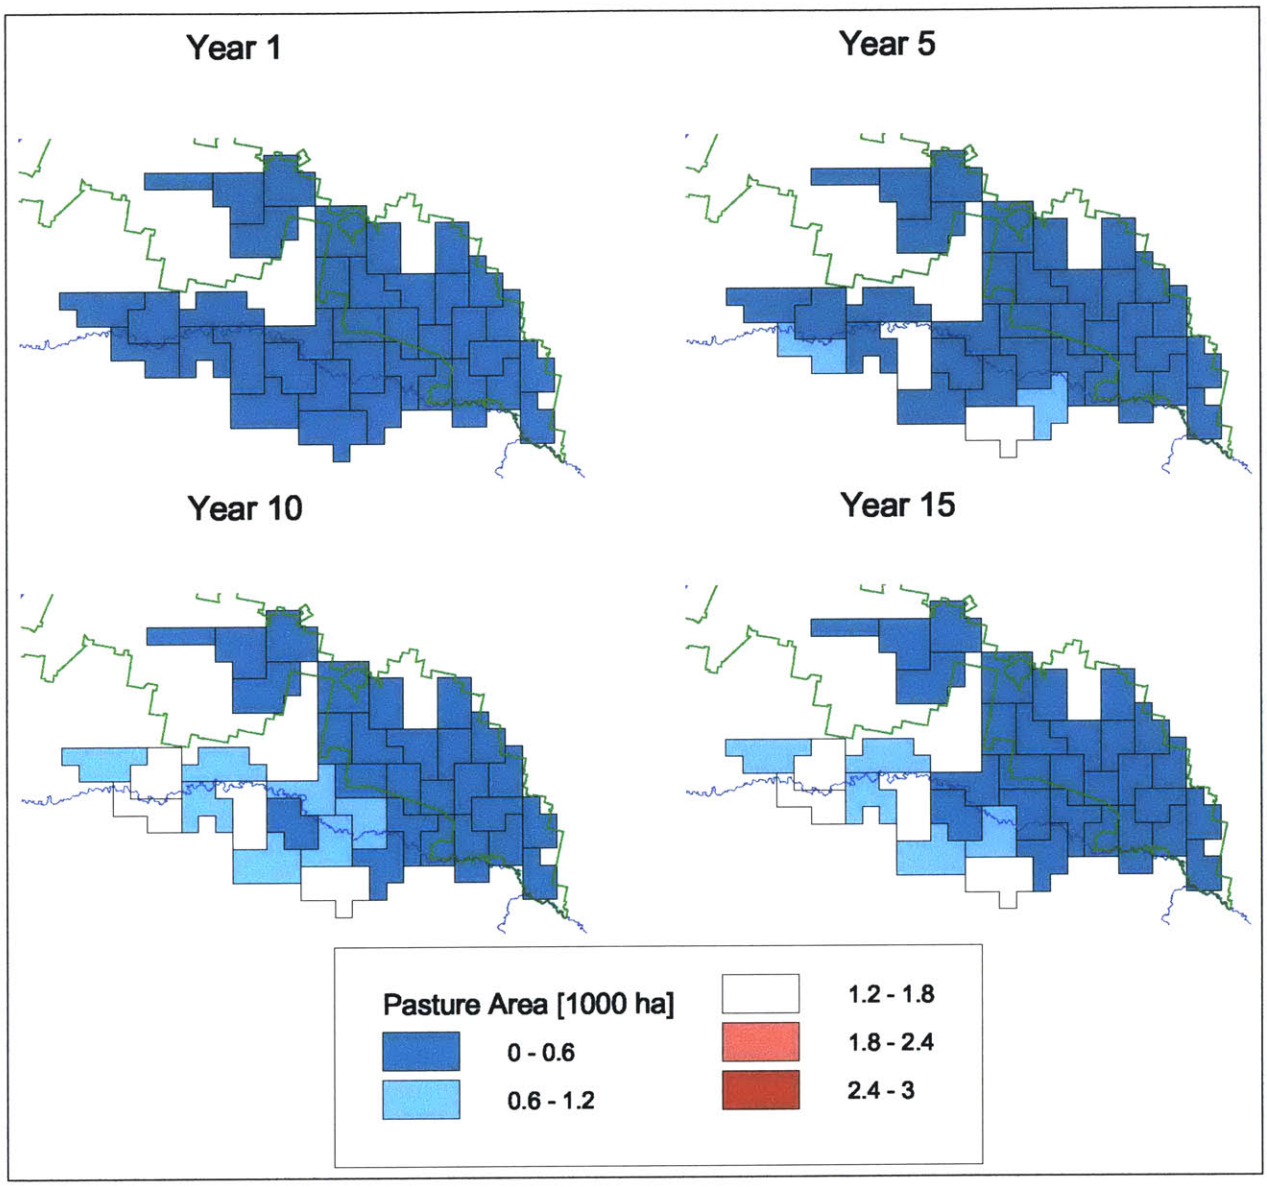

Figure 6.6: Pasture Area Distribution at Selected Times (Scenario 1)

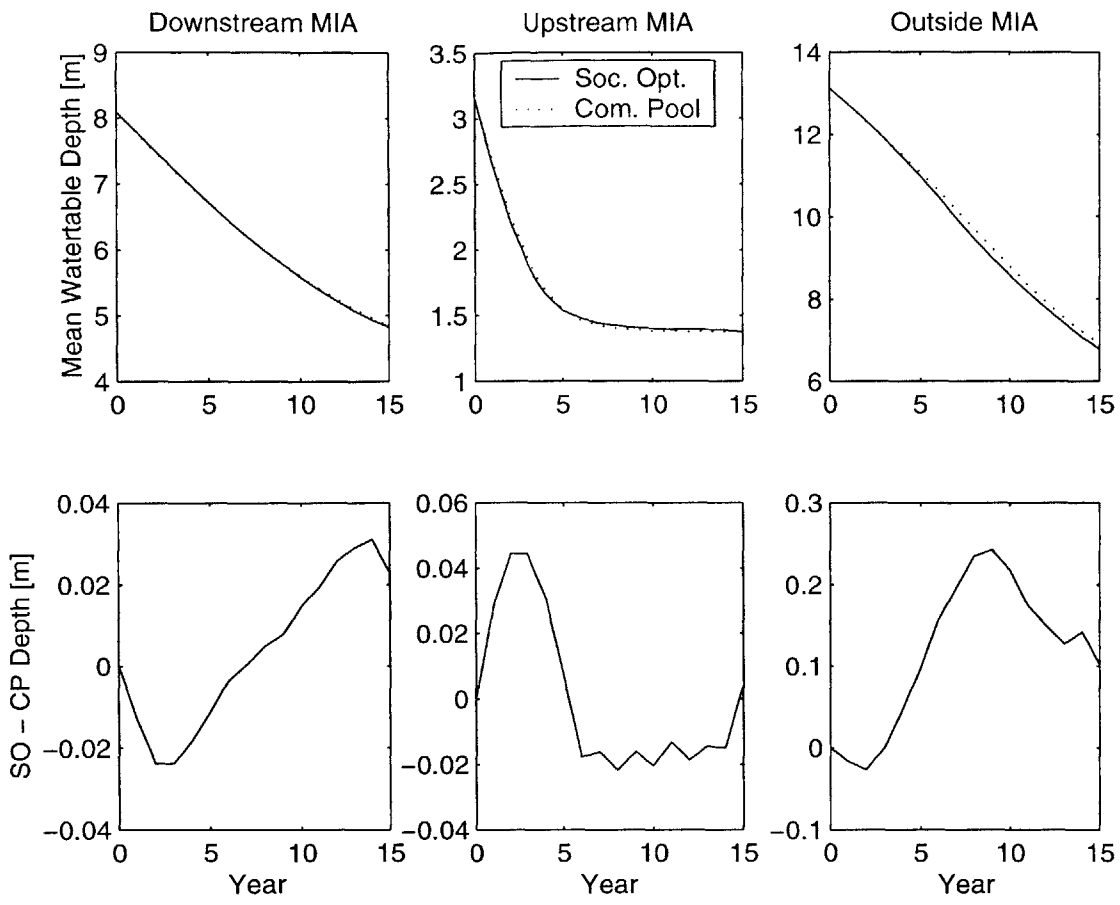

Figure 6.7: Mean Watertable Depth by Zone With Water Market (Scenarios 1 and 2)

We see that there are both positive externalities (shown in red) and negative externalities (shown in blue) resulting from the common pool groundwater arrangement. The negative externality stems from the decrease in yields due to salinization from high watertables. The positive externality stems from the reduction in rice irrigation demand which also results from high watertables. Which effect is largest depends on the current watertable depth and to a lesser degree on the properties of the economic unit. The reduction in recharge effect is felt over a much larger range of depths, so it is dominant when the watertable is deeper than 3 meters. When the watertable is very high, the yield effect becomes more dominant. At the final time, the shadow prices are zero since we have assumed there is no residual or salvage value to the groundwater resource at the end of the planning horizon.

The shadow price of irrigation water is shown in Figure 6.10. In all years, the shadow price is above the actual price of \$11 per ML. This suggests that the current price is too low and that more water is being used for irrigation than is economically efficient. The shadow price is falling over most of the horizon indicating that water is becoming less scarce. More water becomes available as groundwater heads rise and less water per hectare is needed for rice irrigation.

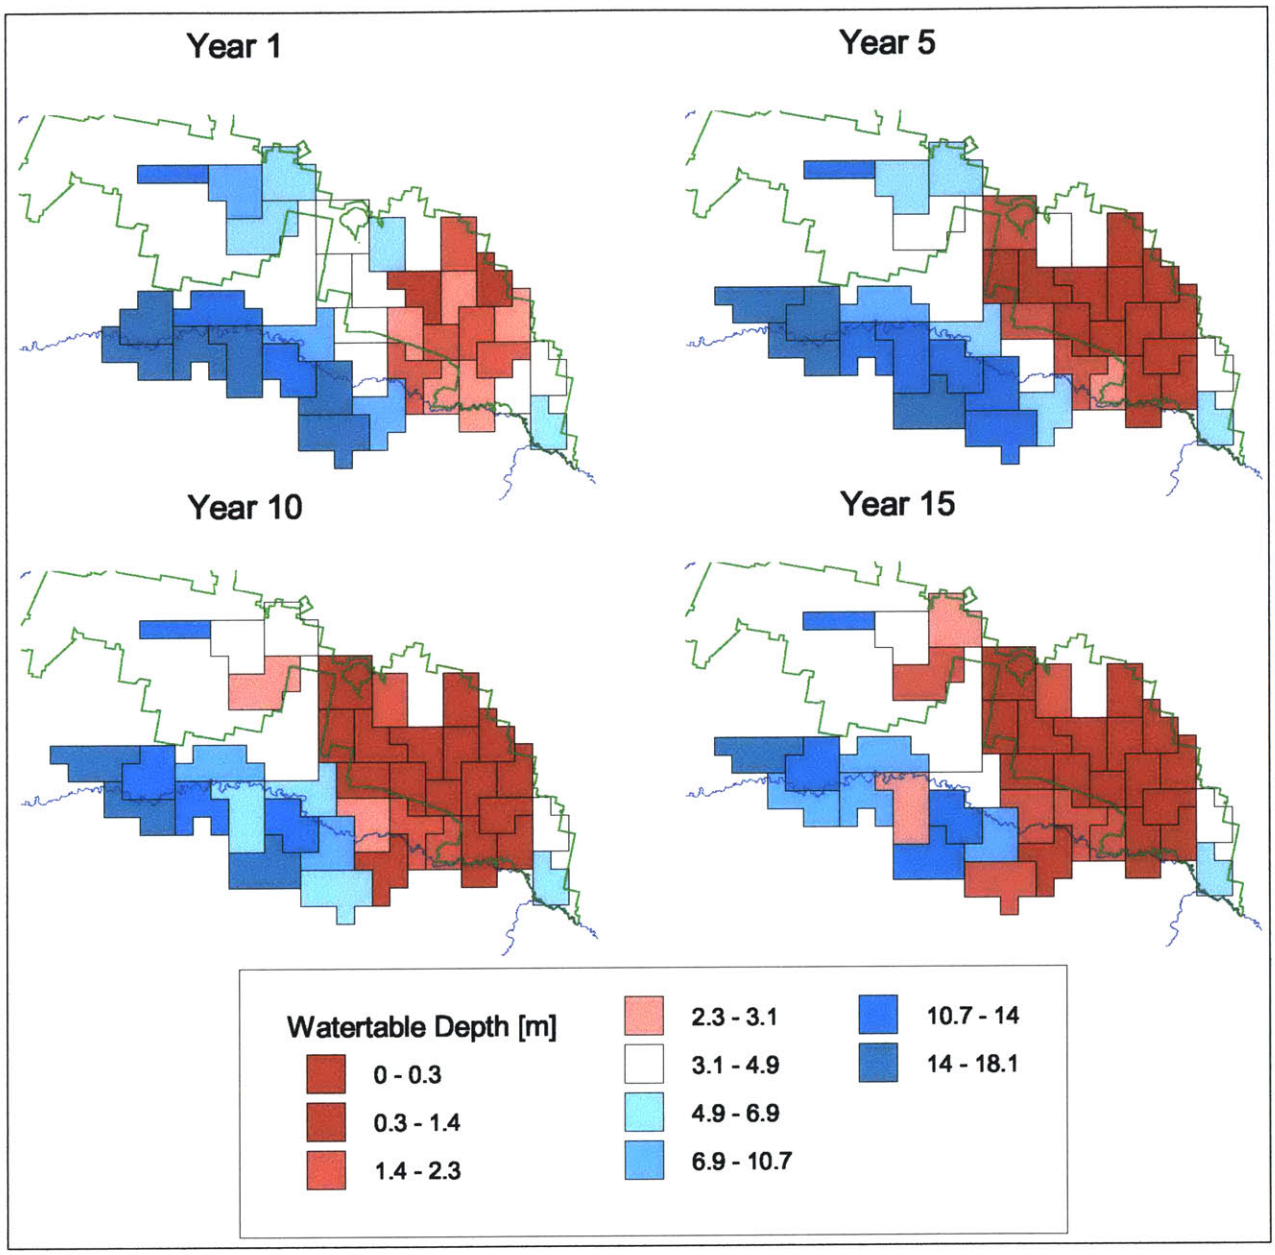

Figure 6.8: Watertable Depth Distribution (Scenario 1)

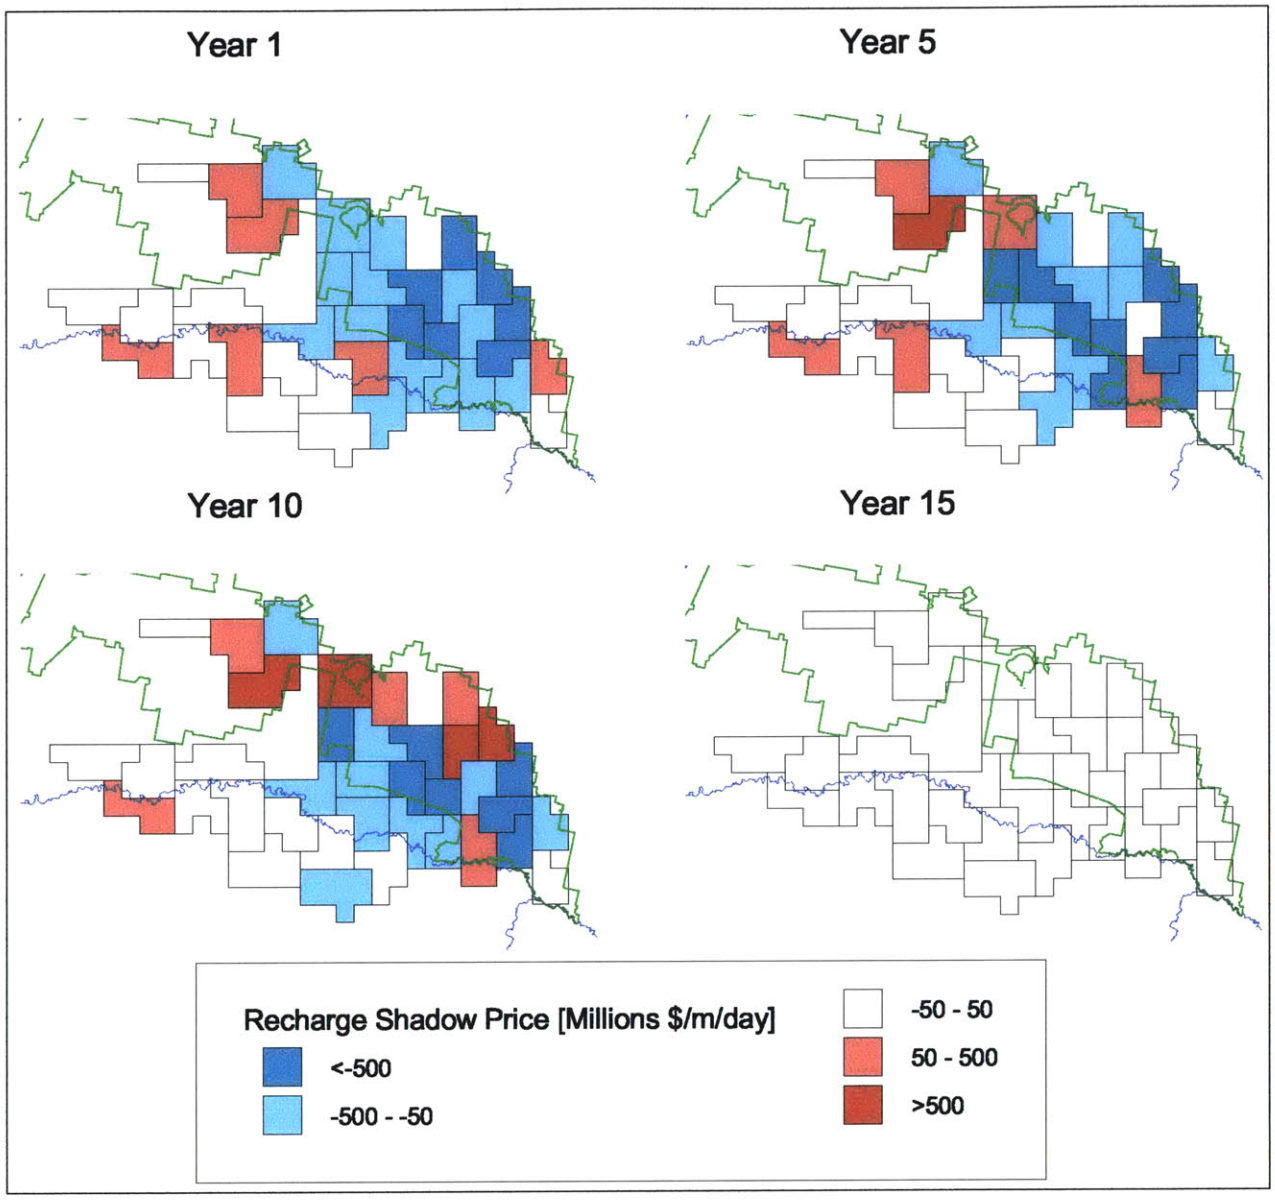

Figure 6.9: Groundwater Recharge Shadow Price Distribution (Scenario 1)

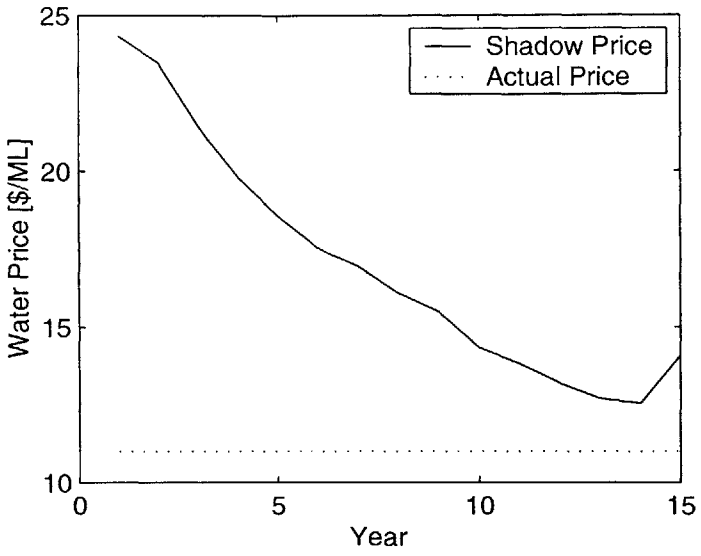

Figure 6.10: Shadow Price of Irrigation Water (Scenario 1)

6.3.2 Results Without Water Market

The results for the scenarios without the quasi-market for water are similar. There is little difference between the social optimum and common pool cases, which means that the cost of the common-pool externality is small. Net benefits are shown in Figure 6.11. Total crop production and crop areas for all economic units are shown Figures 6.12 and 6.13. The spatial distributions of these quantities are very similar to the scenarios with a water market so the spatial plots are not shown.

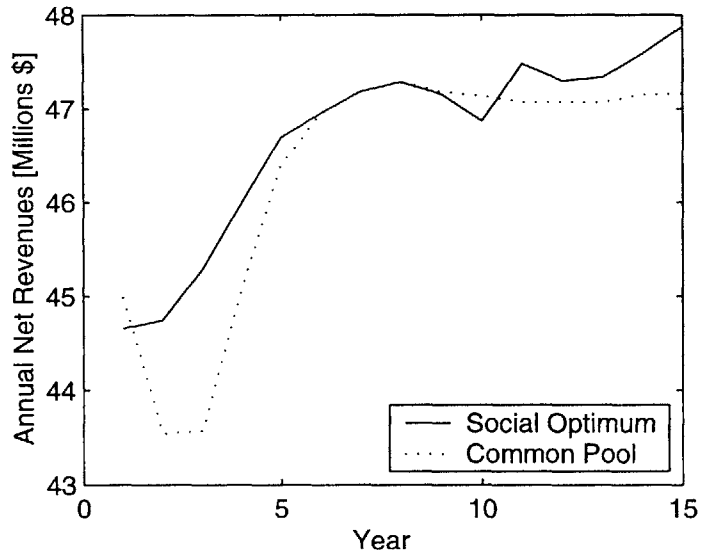

Figure 6.11: Total Net Revenues Without Water Market (Scenario 3 and 4)

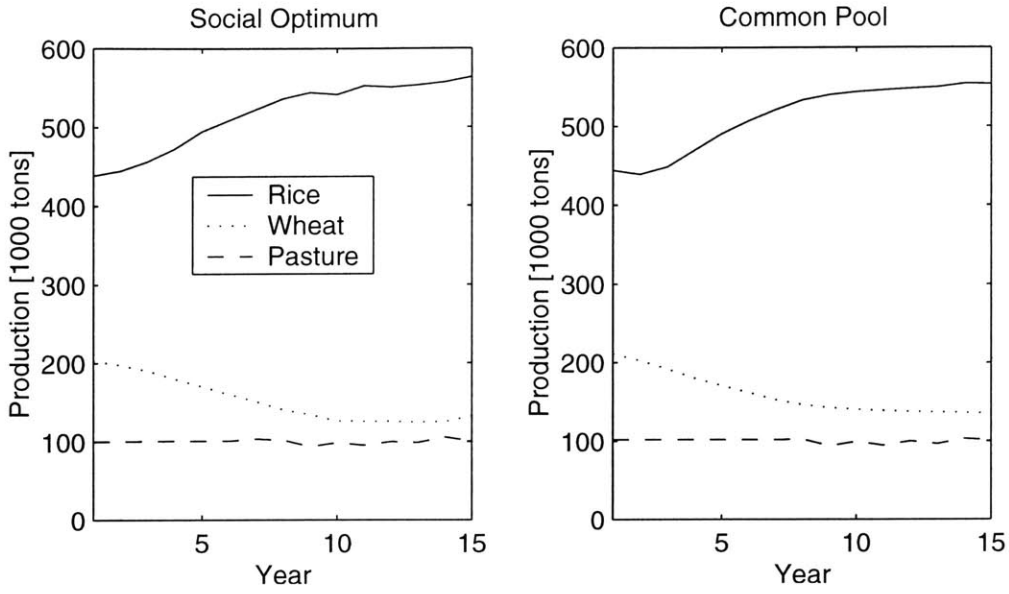

Figure 6.12: Total Crop Production Without Water Market (Scenarios 3 and 4)

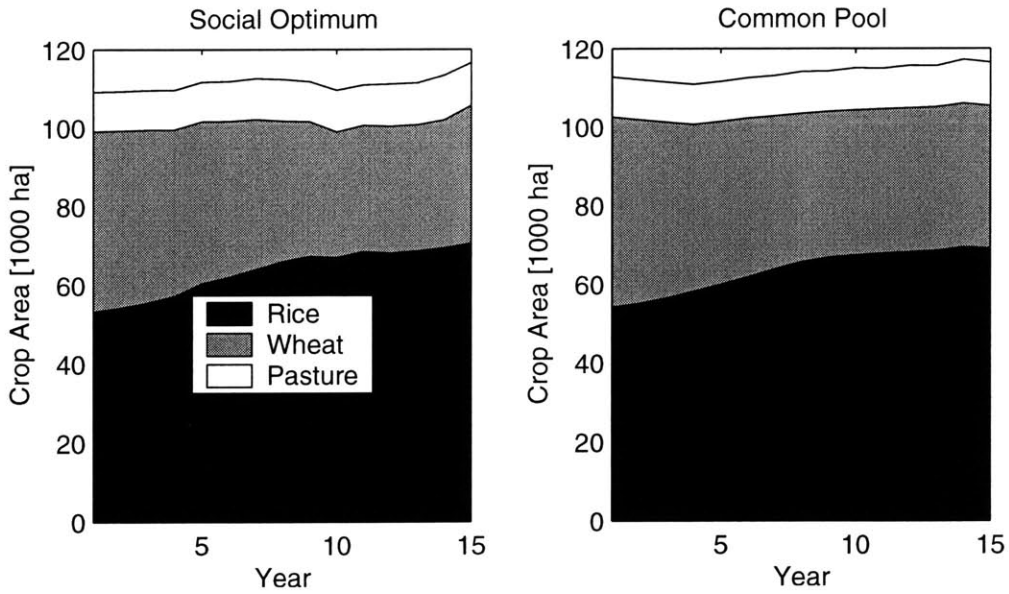

Figure 6.13: Total Crop Areas Without Water Market (Scenarios 3 and 4)

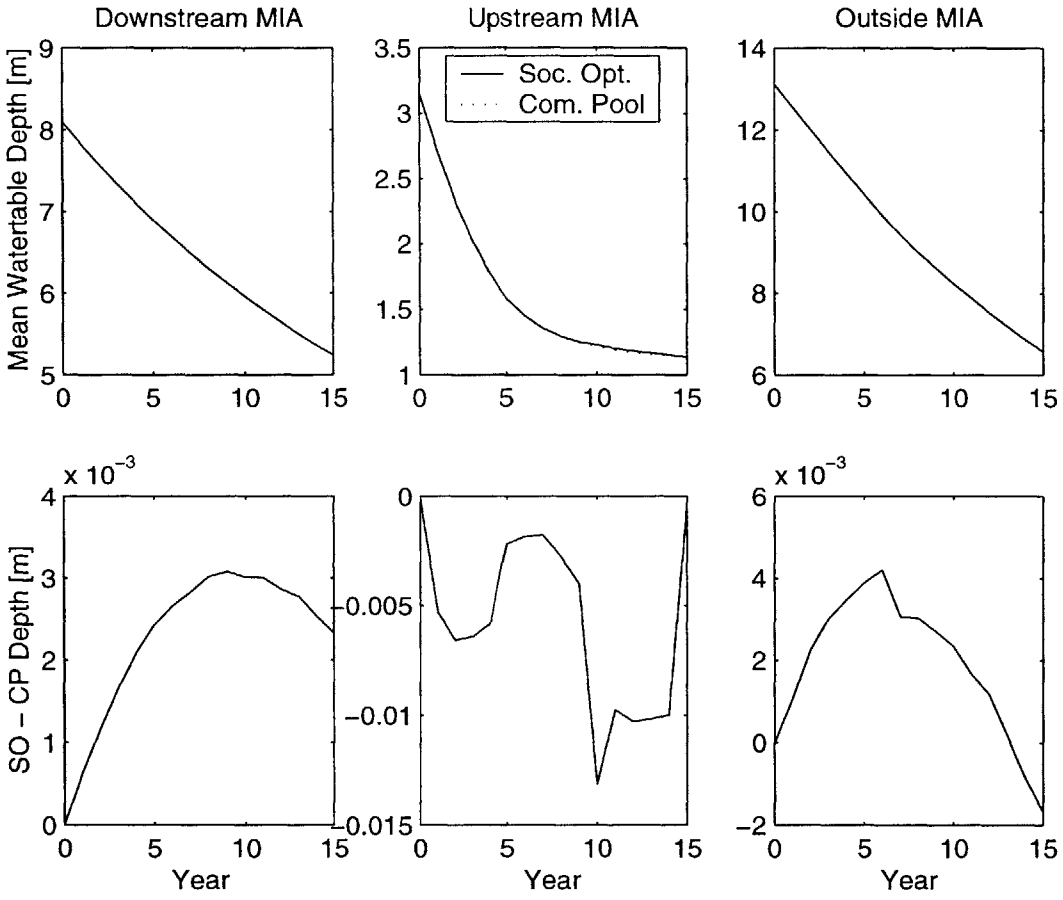

Figure 6.14: Mean Watertable Depth by Zone Without Water Market (Scenarios 3 and 4)

6.4 Common-Pool Policy Option Scenarios

In this section, we present the results of the 30-year common pool scenarios under different combinations of the two policy options. By examining the total discounted revenues of simulations with and without each policy option, we can determine the benefits to society of implementing it. We first show the results of all four scenarios and then analyze the benefits of the rice restriction and the water market.

The annual net revenues of each of the four policy option scenarios are shown in Figure 6.15. The figure also gives the total discounted net revenues in the legend. Total annual crop production and crop areas for the policy scenarios are shown in Figures 6.16 and 6.17. Watertable depths averaged over all of the economic units are shown in Figure 6.18.

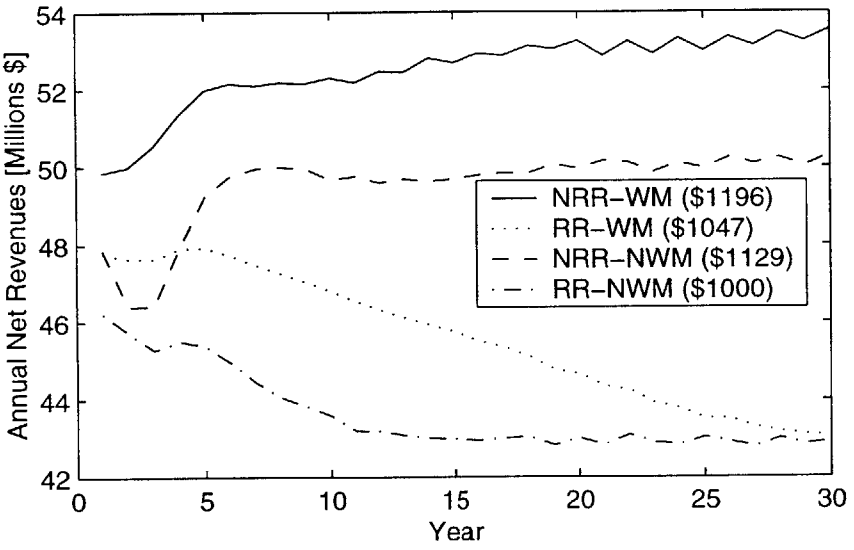

Figure 6.15: Total Annual Net Revenues of Policy Option Scenarios 5–8 In the legend, RR = rice restriction, NRR = no rice restriction, WM = water market, NWM = no water market. The numbers in parentheses are the total discounted net revenues.

6.4.1 Benefits of Rice Restriction

In Section 6.3, we examined the cost of the common pool externality predicted by our modeling framework. The estimated cost was not very large suggesting that it may not be beneficial to regulate land or water use in order to correct it. Given the assumptions of our model, the current restriction on the rice areas may in fact lead to a reduction rather than an increase in net benefits. In this section, we quantify the cost of maintaining this rice area restriction over a 30 year horizon. There are two cases to consider, the case with a water market and the case without a water market.

When there is a water market, we see from Figure 6.15 that the total discounted net revenue without the rice restriction is $1,196 million and with the rice restriction is $1,047

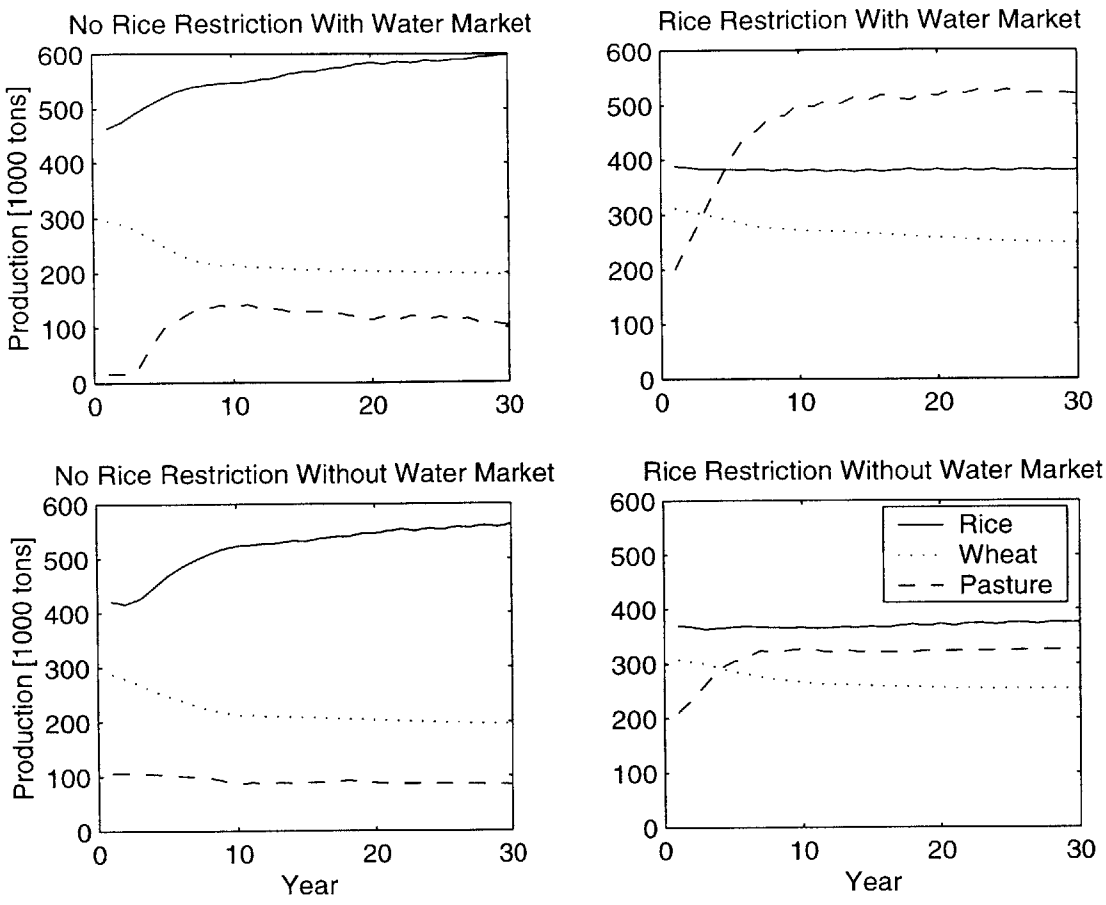

Figure 6.16: Total Annual Crop Production of Policy Option Scenarios 5–8

million. The difference of $149 million is the cost over 30 years of maintaining the rice area constraint. This is a decrease in revenue of 12%. The results are similar when there is no water market. The total discounted net revenue without the rice restriction is $1,129 million and with the restriction is $1,000 million. The cost of maintaining the restriction over 30 years is $129 million, equivalent to an 11% reduction in revenues.

The rice restriction is clearly a binding constraint on production levels. This can be seen in both the production levels shown in Figure 6.16 and the crop areas shown in Figure 6.17. We do not consider potential marketing and processing capacity constraints on rice production. The results indicate, however, a significant economic incentive to increase rice production.

As shown in Figure 6.18, the rice area restriction does maintain deeper watertables but the effect is fairly small. With the water market in place, the rice restriction maintains a greater average watertable depth until year 28. At this point the depth with the rice restriction becomes smaller than the depth without the restriction. The maximum difference in average depths occurs around year 21.

The rice restriction is more effective in the case without the water market. The average watertable depth is always greater with the rice restriction. The maximum different occurs

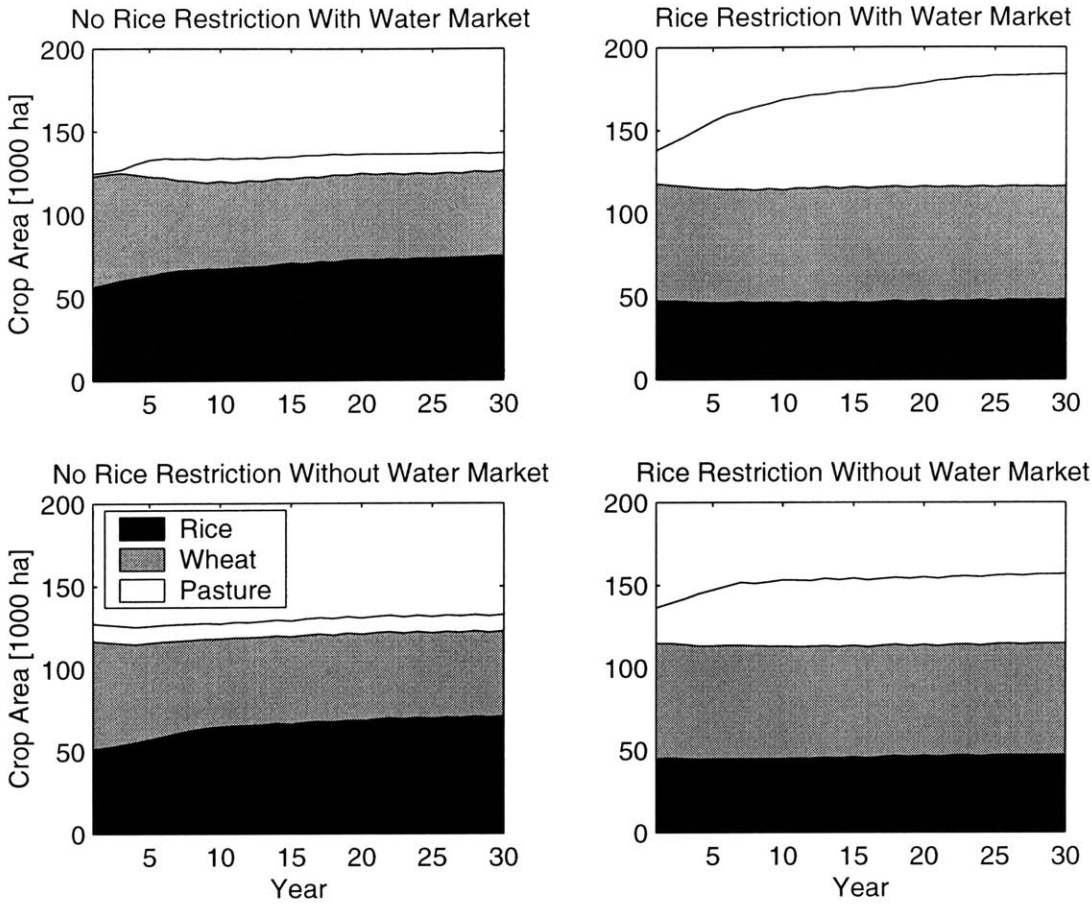

Figure 6.17: Total Annual Crop Areas of Policy Option Scenarios 5–8

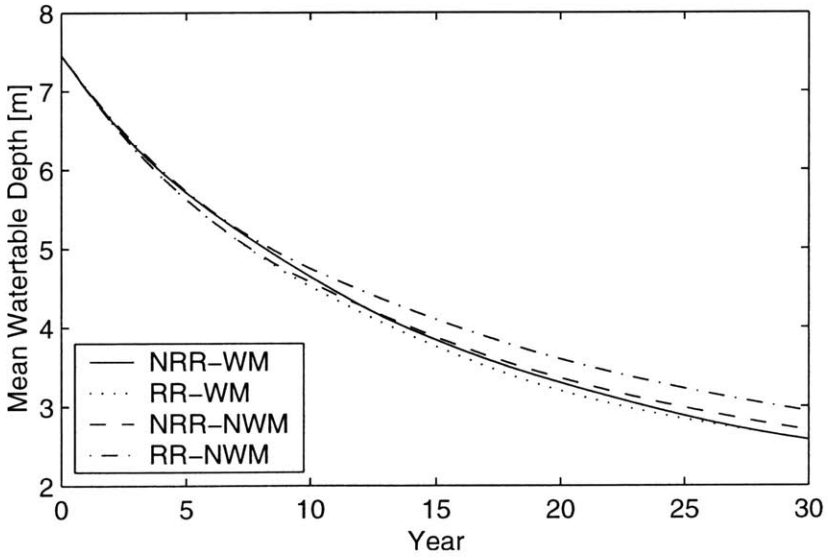

Figure 6.18: Mean Groundwater Depth of Policy Option Scenarios 5–8

near the final time, and is around one-fourth of a meter. In both the water market and no water market cases, farmers adjust to the rice restriction by growing more of the other crops. The groundwater recharge produced by these crops is almost the same as the recharge from the rice crops they replace making the effectiveness of the rice restriction minimal.

6.4.2 Benefits of Water Market

In contrast to the policy restricting rice areas, the water market appears to provide a net benefit to the region. We can estimate these benefits in the case with the rice restriction and the case without it. Starting with the case of no rice restriction, from Figure 6.15, the total revenue with the water market is \$1,196 and without the water market is \$1,129. The difference of \$67 million is the approximate net benefit of implementing the market. Similarly, the benefit in the case of the rice restriction is \$47 million. Since we have modeled the water market in a very simple way, these benefits are only rough approximations. In particular, we have ignored the transaction costs of trading water. Taking these into account would reduce the net benefits from trading. However, the fact that fairly extensive trading is currently taking place suggests that the market is actually providing a net benefit.

The presence of a water market does not significantly change the crop mix when there is no rice restriction. It only leads to a small increase in rice production. This can seen from the plots in the left-hand column of Figures 6.16 and 6.17. When there is a rice restriction, the water market leads to the production of significantly more pasture. This is shown in the plots in the right-hand column of the same figures. When there is no water market, not all of the available water is used. In some economic units applying all of the allotted water would cause the watertable to rise too far. Since the water cannot be transferred to another unit, it goes unused.

This reduction in total water use in the absence of the water market also explains the difference in the average watertable depths in Figure 6.18. The water market allows water to be transferred to areas where rice production could be profitably expanded. This leads to an increase in deep percolation and consequently the watertable rises faster with a water market than without.

6.4.3 Spatial Distribution of Watertable Depth and Crop Areas

We show the spatial distribution of the watertable depth and crop areas for the case with the rice restriction and without a water market (Scenario 6). The distributions for the other scenarios are fairly similar. The watertable depth is shown in Figure 6.19. Crop areas are shown in Figures 6.20, 6.21 and 6.22.

6.5 Long-term Response

In this final section, we investigate the long-term behavior of the groundwater system. Scenarios 9 and 10 are identical to Scenarios 6 and 7 except they are run for 150 years. Scenario 9 has a water market and no rice restriction and Scenario 10 has no water market but does

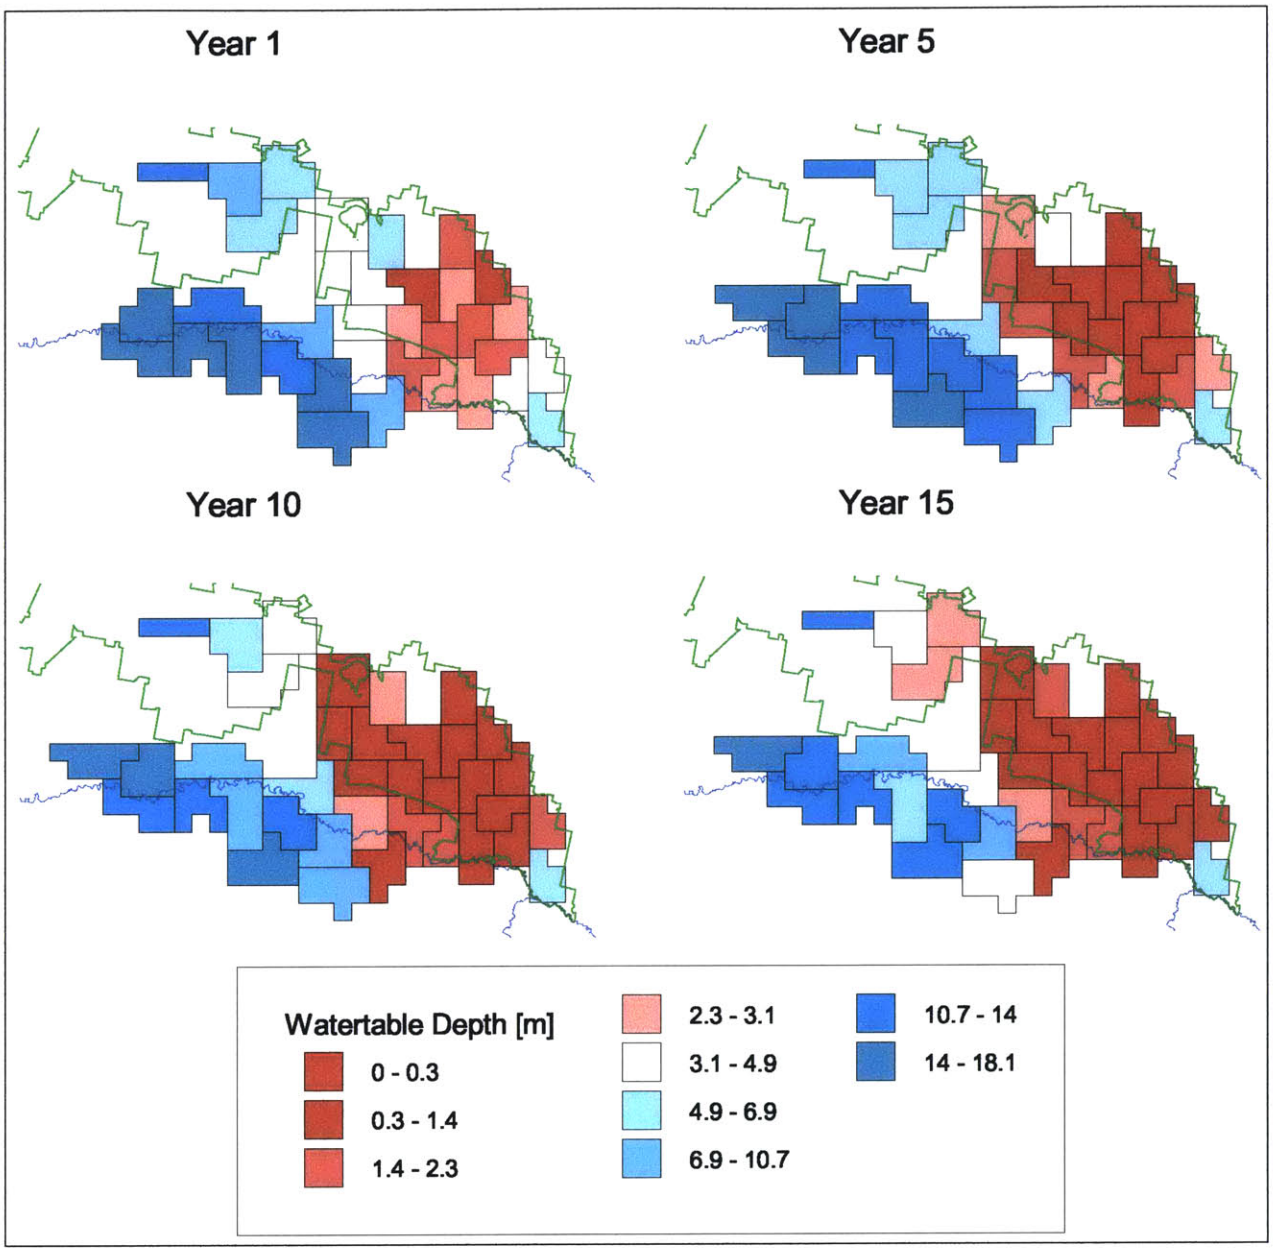

Figure 6.19: Watertable Depth Distribution (Scenario 6)

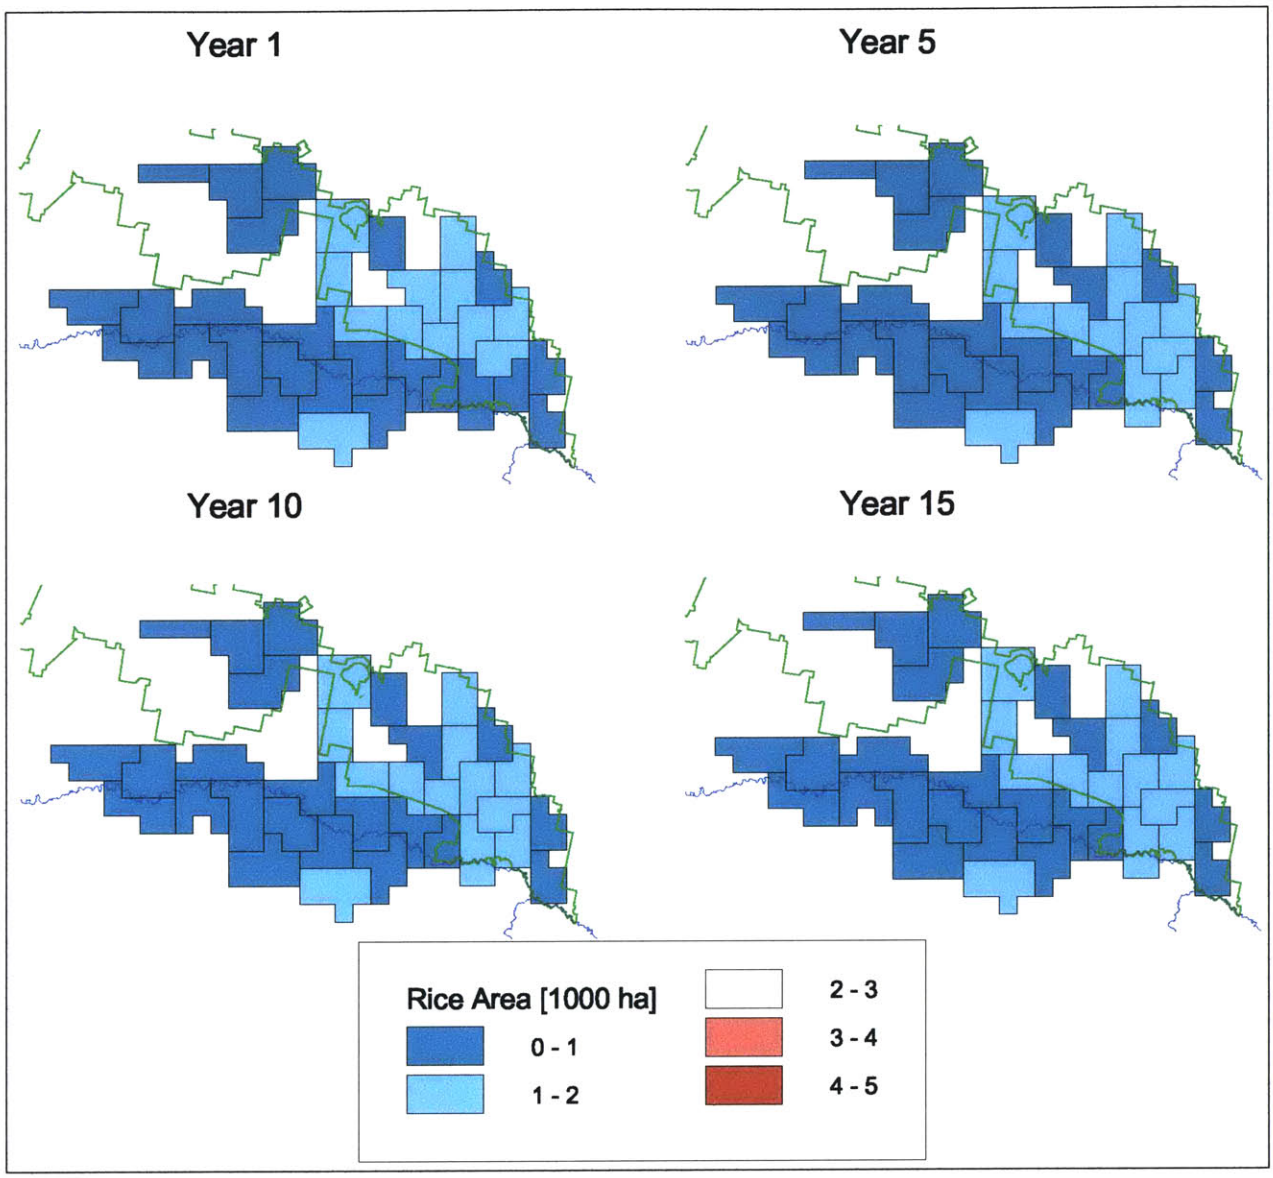

Figure 6.20: Rice Area Distribution at Selected Times (Scenario 6)

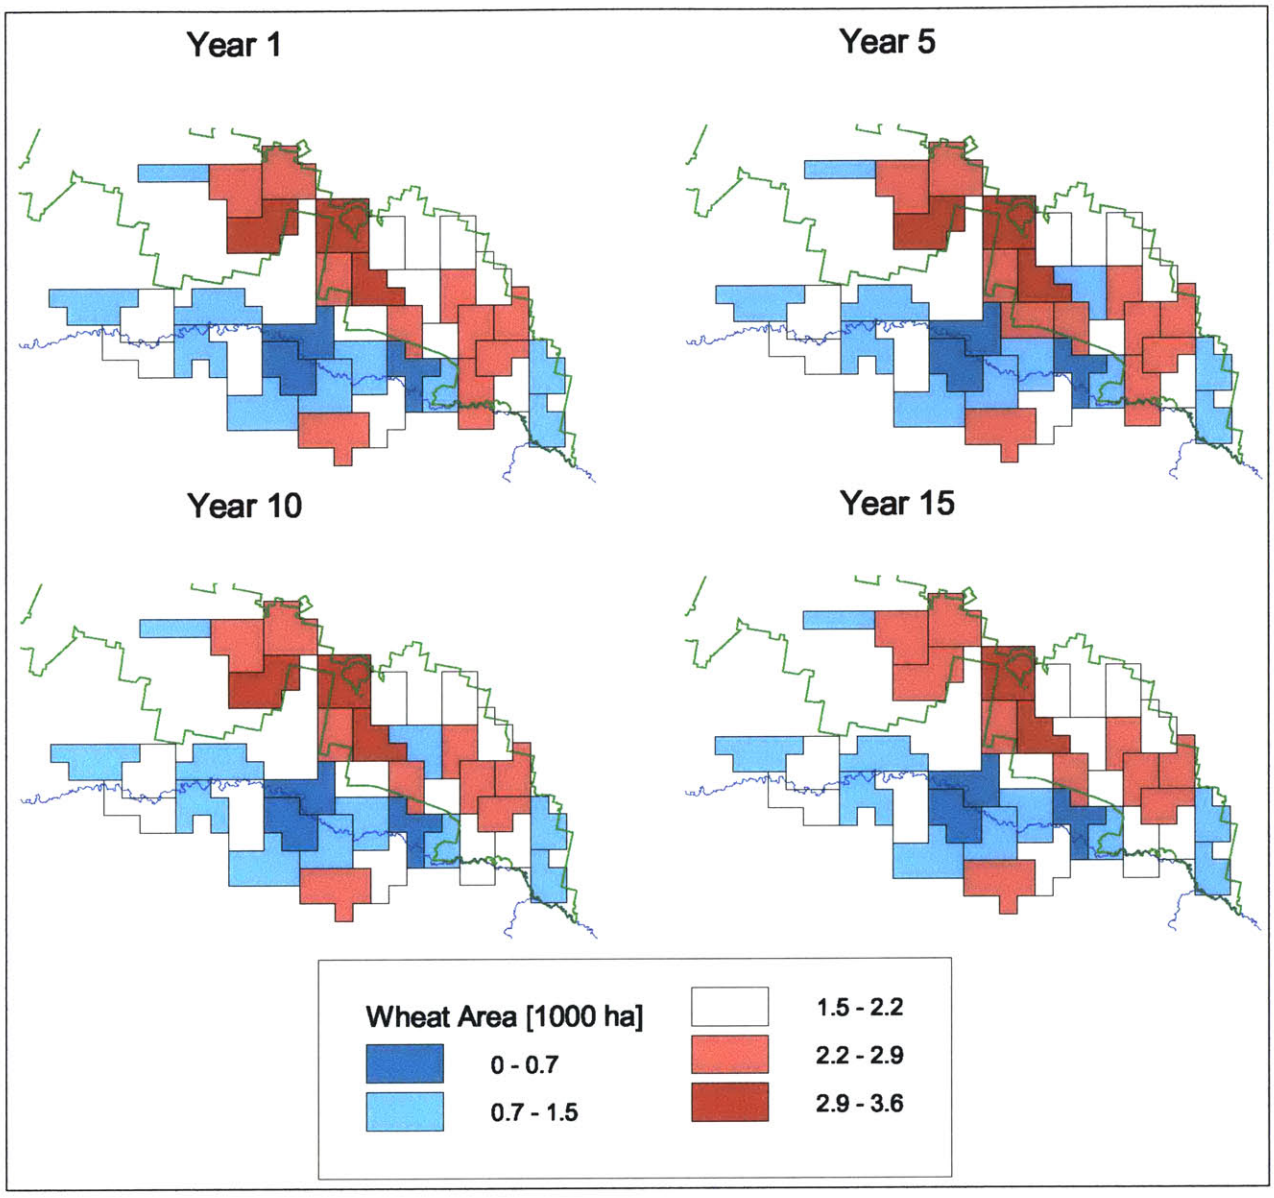

Figure 6.21: Wheat Area Distribution at Selected Times (Scenario 6)

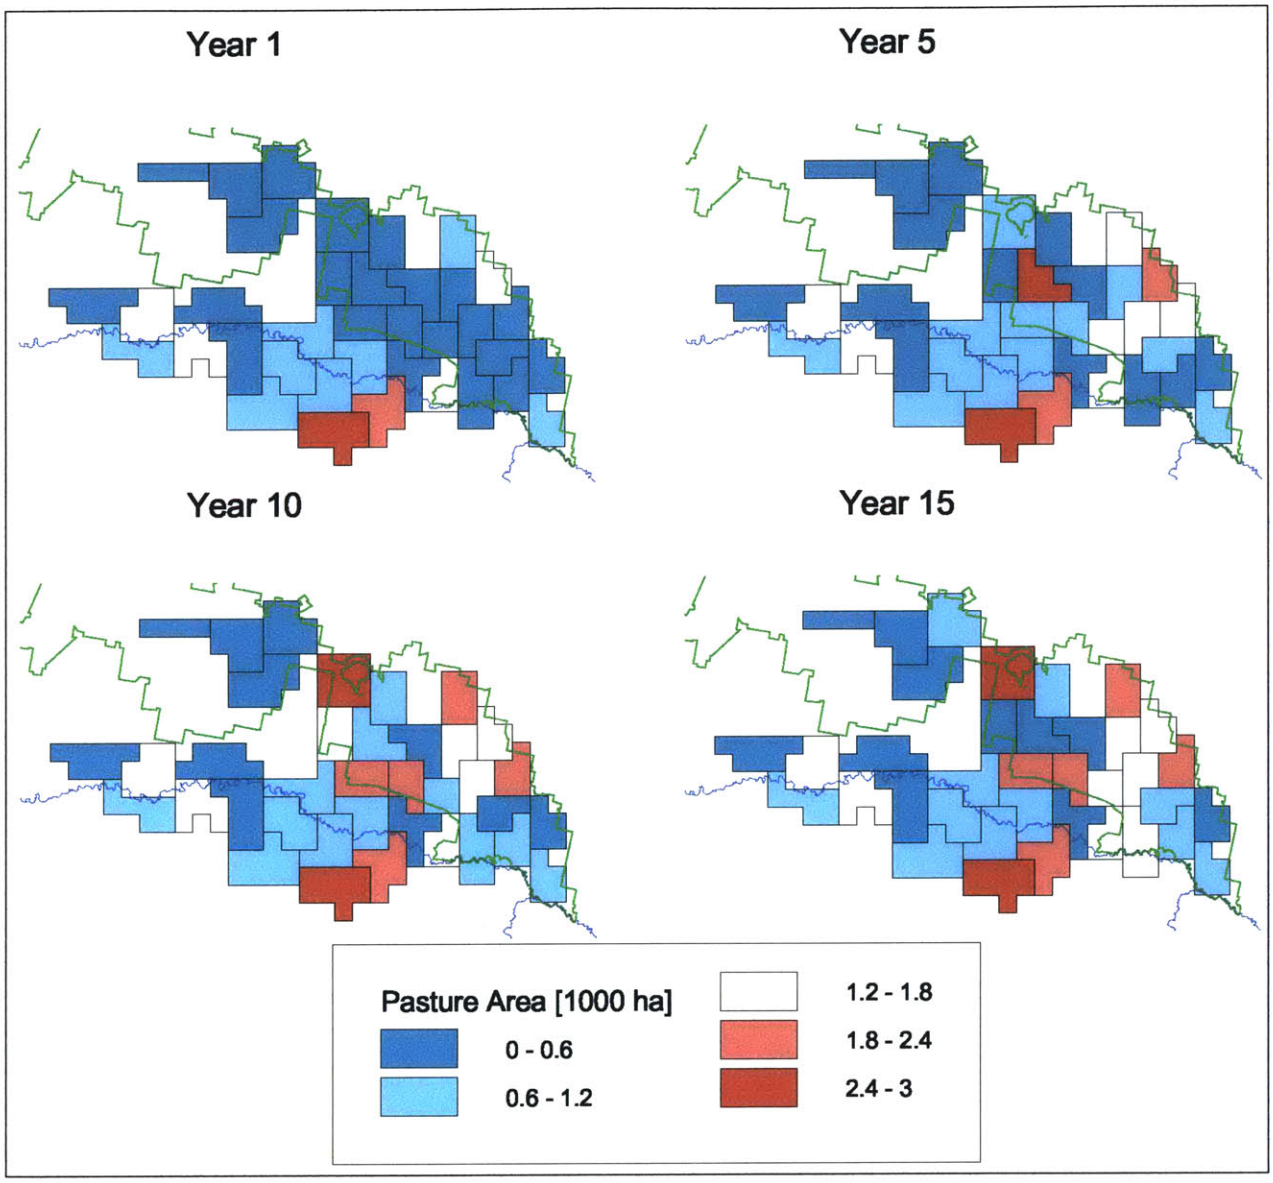

Figure 6.22: Pasture Area Distribution at Selected Times (Scenario 6)

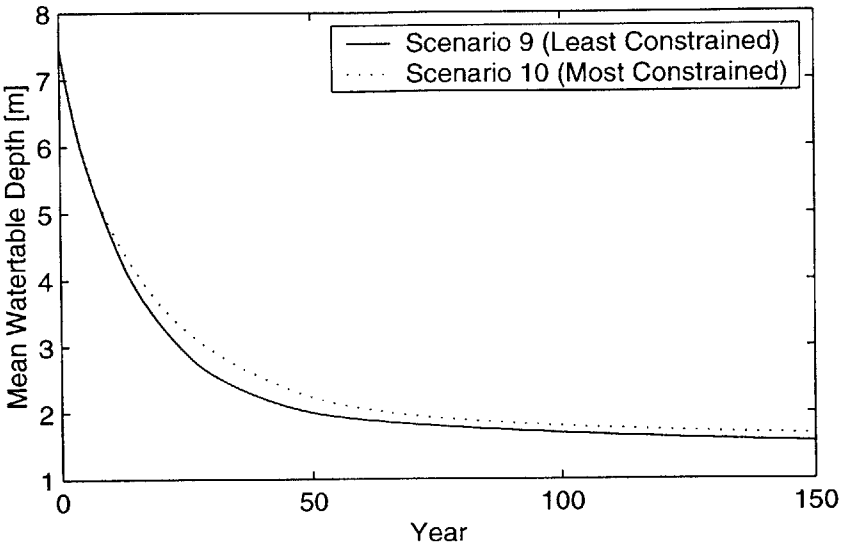

Figure 6.23: Long-term Mean Groundwater Depth (Scenarios 9–10)

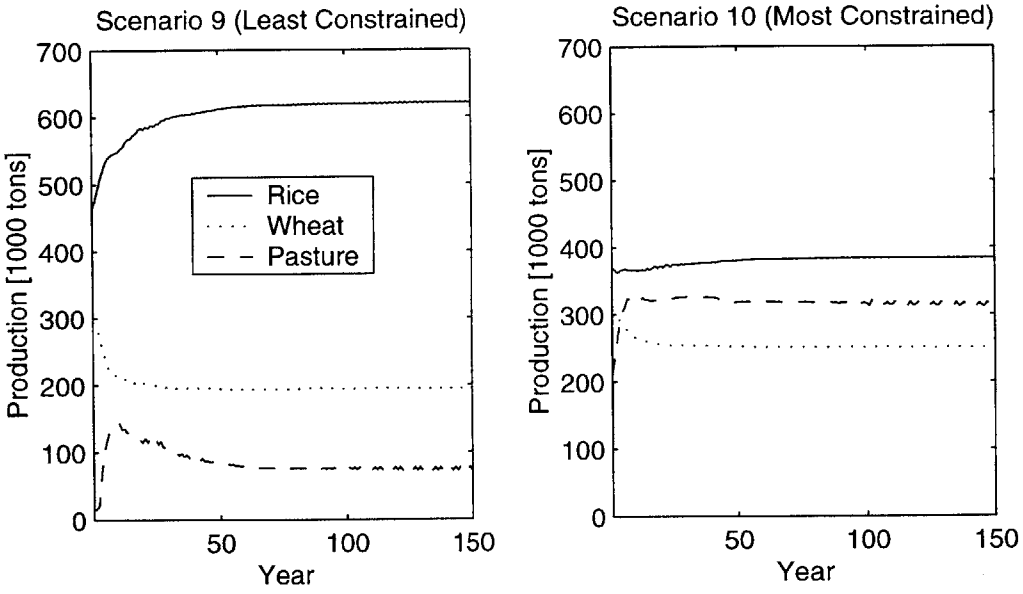

Figure 6.24: Total Annual Crop Production (Scenarios 9–10)

have the rice restriction. These were chosen because they represent least constrained and the most constrained cases and so will provide the biggest contrast.

The mean groundwater depths are shown in Figure 6.23, and the corresponding levels of crop production are shown in Figure 6.24. After around 100 years the system is close to a steady-state, although the depth to groundwater is still slowly decreasing. The final mean watertable elevation is very high, just under 2 m below the ground surface. For both scenarios, this situation appears sustainable since crop production continues at reasonable steady-state levels.

There are other considerations which are not captured in our modeling framework which could also be important to sustainability. These include both questions of equity within the study area and potential effects on conditions outside of the study area. In addition, the large space and time scales of our analysis necessarily average out important small-scale processes. Some of these limitations are noted in the next chapter as areas for future research.

Chapter 7

Conclusions

In this final chapter, we provide a summary of the contributions of this research and areas for future work.

7.1 Summary of Original Contributions

The research described in this thesis has made original contributions in three main areas. We have: 1) developed a hydrologic-economic modeling framework which can be used for policy analysis of regional salinization issues; 2) demonstrated how a model order reduction technique can be used to reduce the computational cost representing groundwater flow in an optimization problem; and 3) applied our integrated modeling framework to a case study of land and water management policies in the Lower Murrumbidgee Catchment of Australia.

7.1.1 Hydrologic-Economic Model of Salinization

We formulated a dynamic and spatially-distributed model of the hydrology and economics of irrigation-induced salinization. The model is unique because it provides an realistic representation of hydrologic processes in an economic framework that is both dynamic and spatially-distributed. It is designed to provide regional-scale analyses of land and water management policies relating to salinization. The model can be solved using standard nonlinear programming software with a modest computational cost.

The model uses a simplified representation of unsaturated zone dynamics derived from a detailed physically-based model. This approach captures the complex dynamics of the unsaturated zone while maintaining computational feasibility.

7.1.2 Model Order Reduction of Groundwater System

Groundwater flow was represented in a computationally efficient way by using the model reduction algorithm of balanced truncation. We are not aware of the algorithm being used

in this application previously. It provides an viable alternative to the response matrix and related methods.

7.1.3 Policy Analysis of Salinization in Murrumbidgee Catchment

As part of the case study of the Murrumbidgee Catchment, we estimated the cost of the groundwater common-pool externality and the benefits of the rice area restriction and the water market.

Cost of Common-Pool Externality

Our results showed that the cost of the externality caused by having a common-pool arrangement for the groundwater system is small. This was true both with and without a simple representation of a water market.

Benefits of Rice Area Restriction and Water Market

Within our modeling framework, the rice restriction brings about negative benefits to the region with or without the presence of a water market. The restriction may still be beneficial if it corrects externalities which are not represented in our model. In particular, the model does not consider externalities relating to risk aversion or meterological and hydrological uncertainty. There may also be externalities which are not captured because they act over scales smaller than the model resolves. We find that the water market yields positive benefits independent of the rice restriction. This is expected since there is a successful water market operating currently. The model also predicts that neither of the policy options will significantly affect the watertable depth trajectory over the longer term.

Long-term Response

We examined the response of the system over a 150-year horizon and found it reaches an approximate steady-state after 100 years. Final average groundwater levels are just under two meters below the ground surface.

7.2 Recommendations for Future Research

7.2.1 Uncertainty and Variability

The current framework is entirely deterministic: it assumes all parameters and inputs are known with certainty. In particular, it ignores the considerable uncertainty in soil and aquifer parameters and of future meteorology and surface water availability. If these uncertainties were present in the model, the results would probably show more conservative hedging behavior with the watertable depth. In other words, the optimal watertable position would probably be deeper than in the deterministic model.

7.2.2 Representation of Crop Dynamics and Yield Functions

The representation of crop growth we used was very simple. Future work could include the effect of waterlogging and sensitivity of the plant to salinity at different growth stages. Wheat and pasture irrigation depths should be determined by water deficits instead of being constant. The component of the production function representing decreasing returns to land should be estimated for each economic unit from actual yield and crop mix data.

The yield response to groundwater depth should be verified against field data. In particular, the rice yield function appears unrealistically insensitive to high watertables until the shallow groundwater head is almost at the ground surface.

7.2.3 Representation of Groundwater Flow and Salt Transport

Further work should be done to test the adequacy of our representations of groundwater flow and salt transport. We assumed that there is only one-dimensional vertical flow in the near-surface zone (0–7 m) and two-dimensional horizontal flow in the deeper zone (>7 m). We ignore the dynamics of salinization of the root-zone and potential changes in groundwater salinity. The groundwater model is run on an annual timestep which disregards potentially important seasonal head changes. There may also be important effects occuring on spatial scales smaller than we currently resolve. Significant hydrogeological features, such as prior stream deposits, are not represented in the groundwater mode. Including these features may be necessary to accurately estimate the magnitude of the common-pool externality. Some of these assumptions could be tested with two-dimensional flow and transport simulations of cross sections which include both the near-surface and deeper zones. This simulation would ideally be validated against field measurements from the study area.

7.2.4 Extension of Lower Murrumbidgee Model

Groundwater Pumping as Control Variable

Currently, groundwater pumping is fixed at current rates. The model would be more realistic if pumping rates in each economic unit were control variables along with the crop choice. This would allow the exploration of the interaction between watertable depths and pumping.

Extend Scope of Murrumbidgee Case Study

The current model could be extended to include all of the Lower Murrumbidgee Catchment. Most of the necessary data describing the Coleambally Irrigation Area has already been collected by the CSIRO in Griffith. In addition, the number of crops simulated could be expanded to include other crops grown in the region, such as lucerne, perennial pasture, and soybeans.

Appendix A

Unsaturated Zone Simulation Parameters and Transfer Functions

A.1 Unsaturated Zone Simulation Parameters

Table A.1: Monthly Crop Factors

	Rice	Wheat	Soybean	Perennial Pasture	Annual Pasture	Lucerne	Fallow
Sow Date[a]	15-Oct	15-May	30-Nov	1-Sep	25-Mar	1-Oct	NA
Growing Period[b]	150	200	140	365	230	334	0
Crop Factors[c]							
Jan	1.10	0.20	0.75	0.85	0.20	1.30	0.20
Feb	1.10	0.20	1.05	0.85	0.20	1.30	0.20
Mar	1.00	0.20	1.00	0.85	0.20	1.20	0.20
Apr	0.40	0.30	0.50	0.85	0.40	1.20	0.30
May	0.20	0.40	0.30	0.80	0.60	1.00	0.30
Jun	0.20	0.60	0.40	0.80	0.70	0.65	0.40
Jul	0.20	0.90	0.40	0.80	0.80	0.65	0.40
Aug	0.30	1.05	0.40	0.80	0.80	0.65	0.40
Sep	0.40	1.05	0.30	0.80	0.80	0.90	0.30
Oct	1.00	0.80	0.30	0.85	0.60	1.20	0.20
Nov	1.10	0.50	0.30	0.85	0.40	1.30	0.20
Dec	1.10	0.20	0.45	0.85	0.20	1.30	0.20

Source: Wu et al. [1998]

[a] Approximate day crop is planted.

[b] Approximate length of growing season in days.

[c] Crop factor is multiplied by the reference ET to give potential crop ET.

A.2 Unsaturated Zone Transfer Functions

A.2.1 Rice

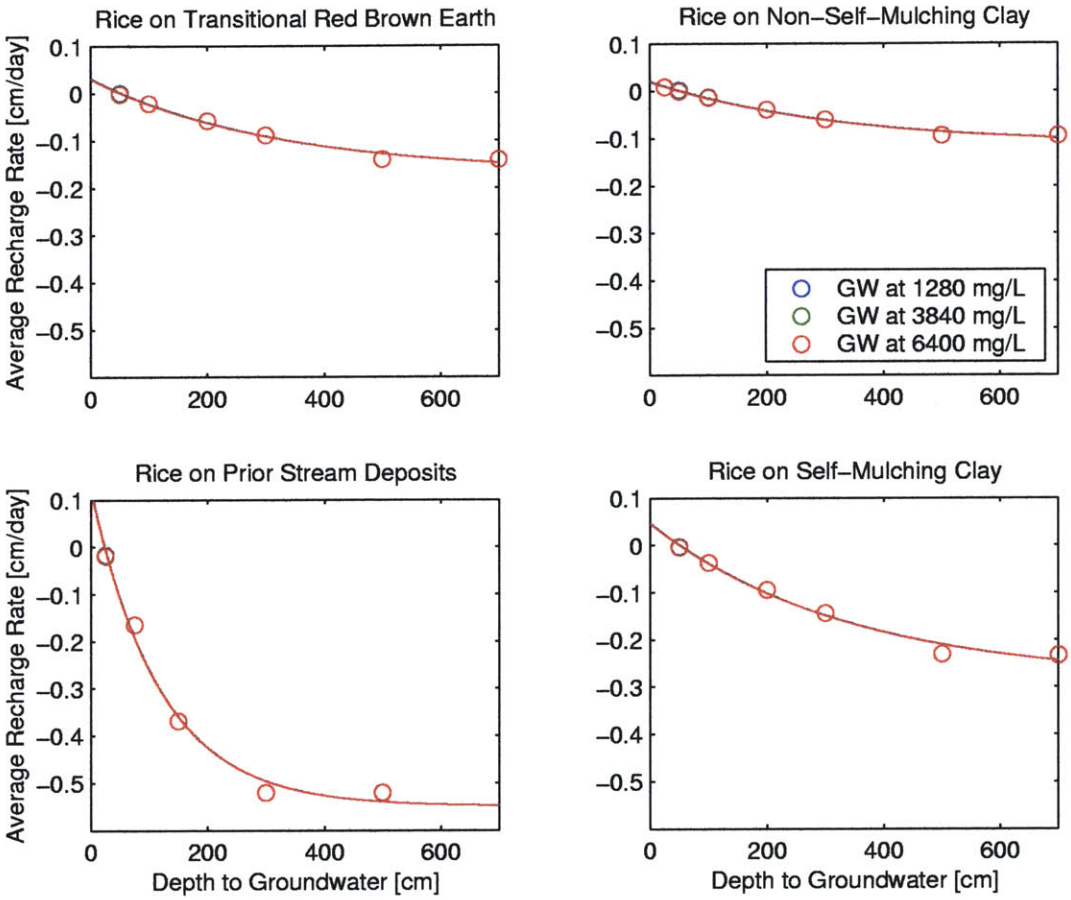

Figure A.1: Rice Recharge Transfer Function

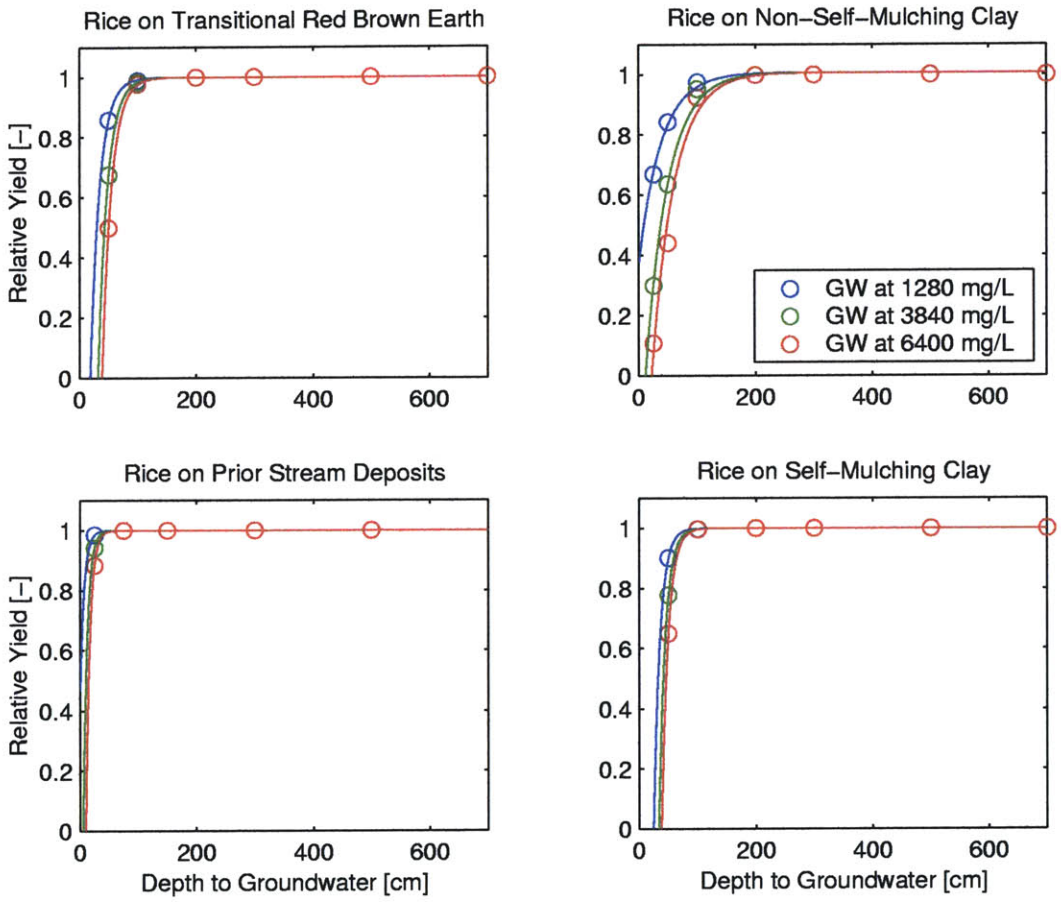

Figure A.2: Rice Relative Yield Transfer Function

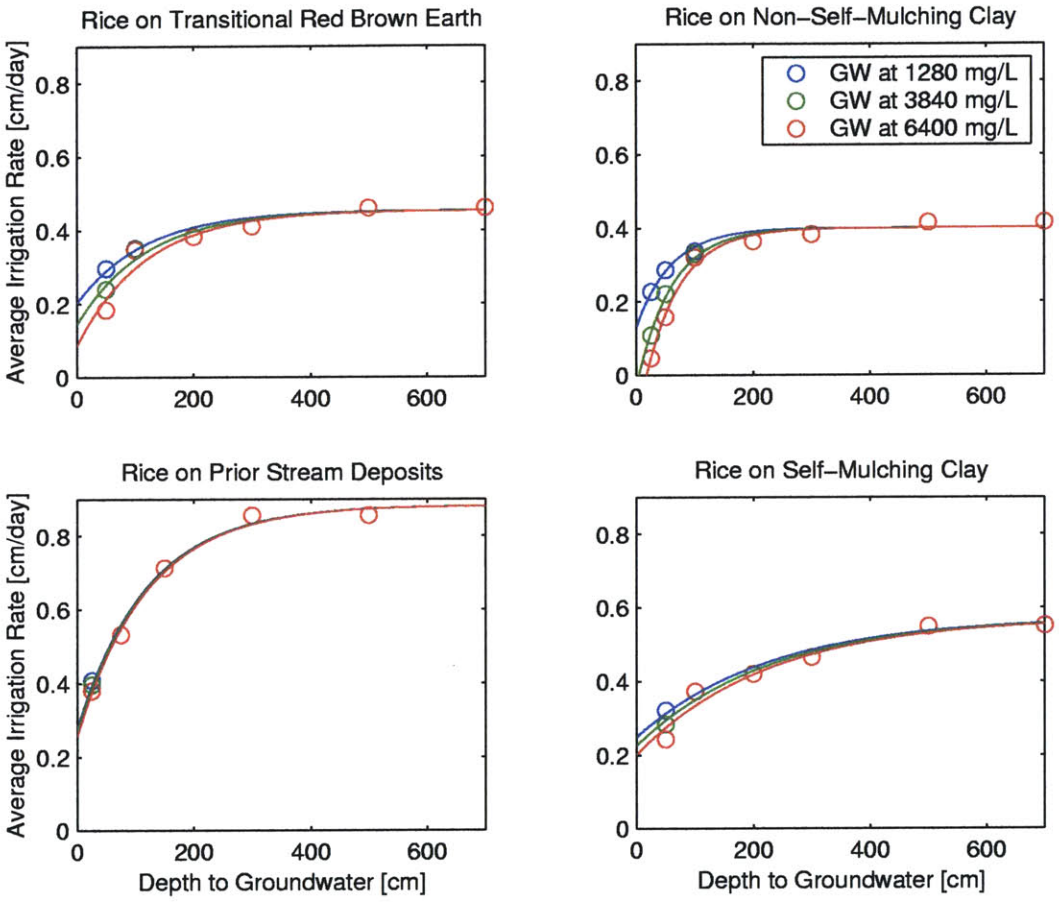

Figure A.3: Rice Irrigation Requirement Transfer Function

A.2.2 Wheat

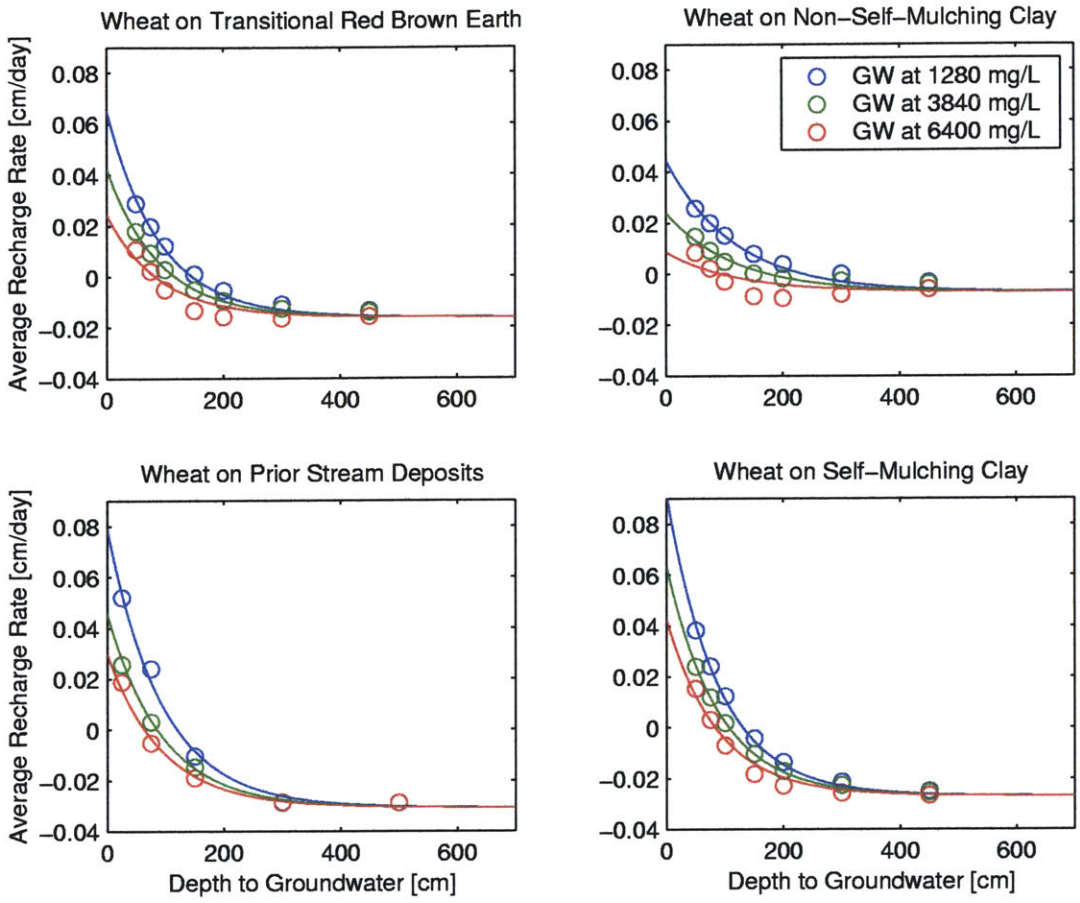

Figure A.4: Wheat Recharge Transfer Function

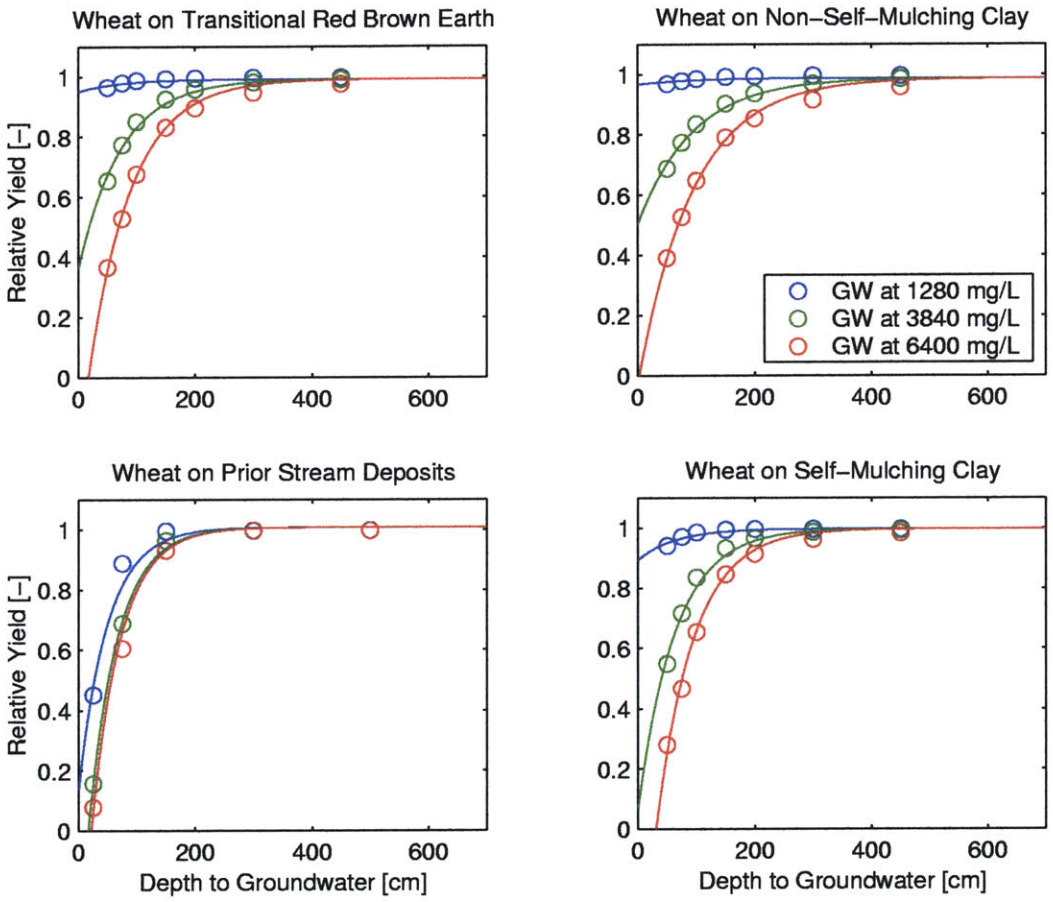

Figure A.5: Wheat Relative Yield Transfer Function

A.2.3 Pasture

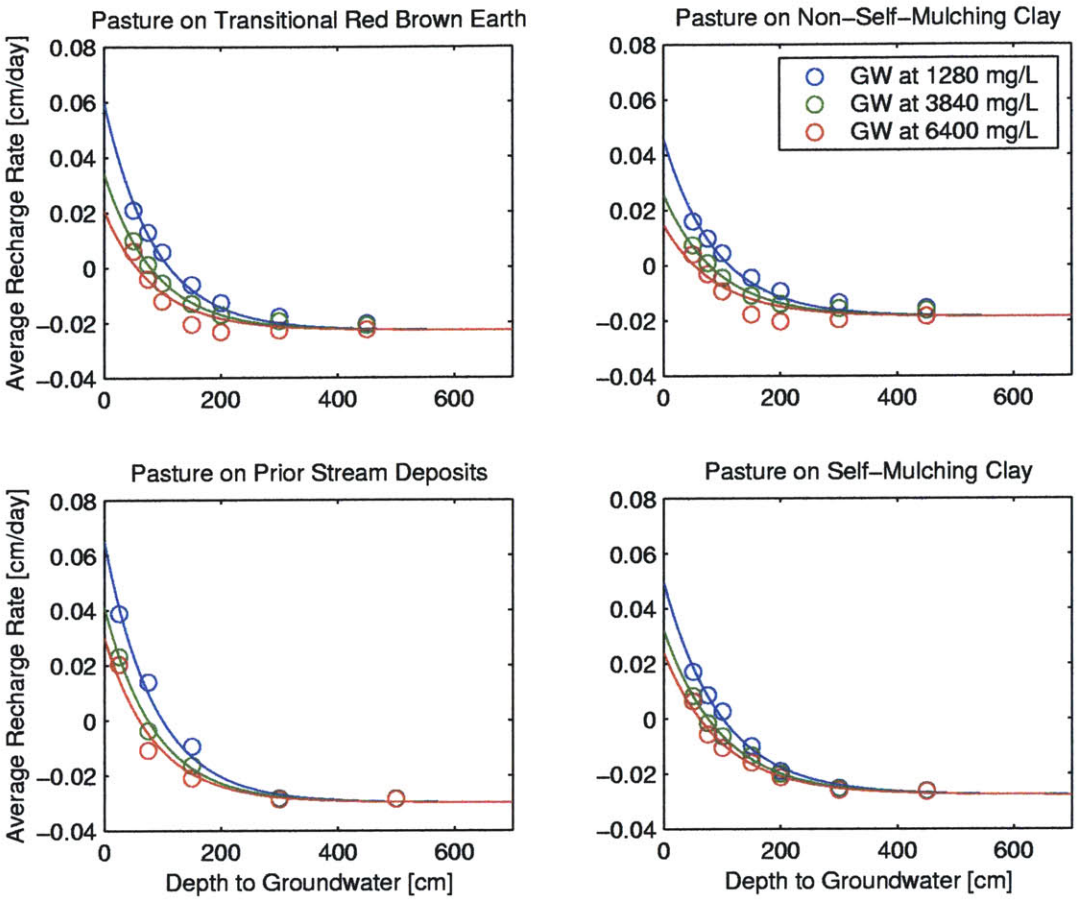

Figure A.6: Pasture Recharge Transfer Function

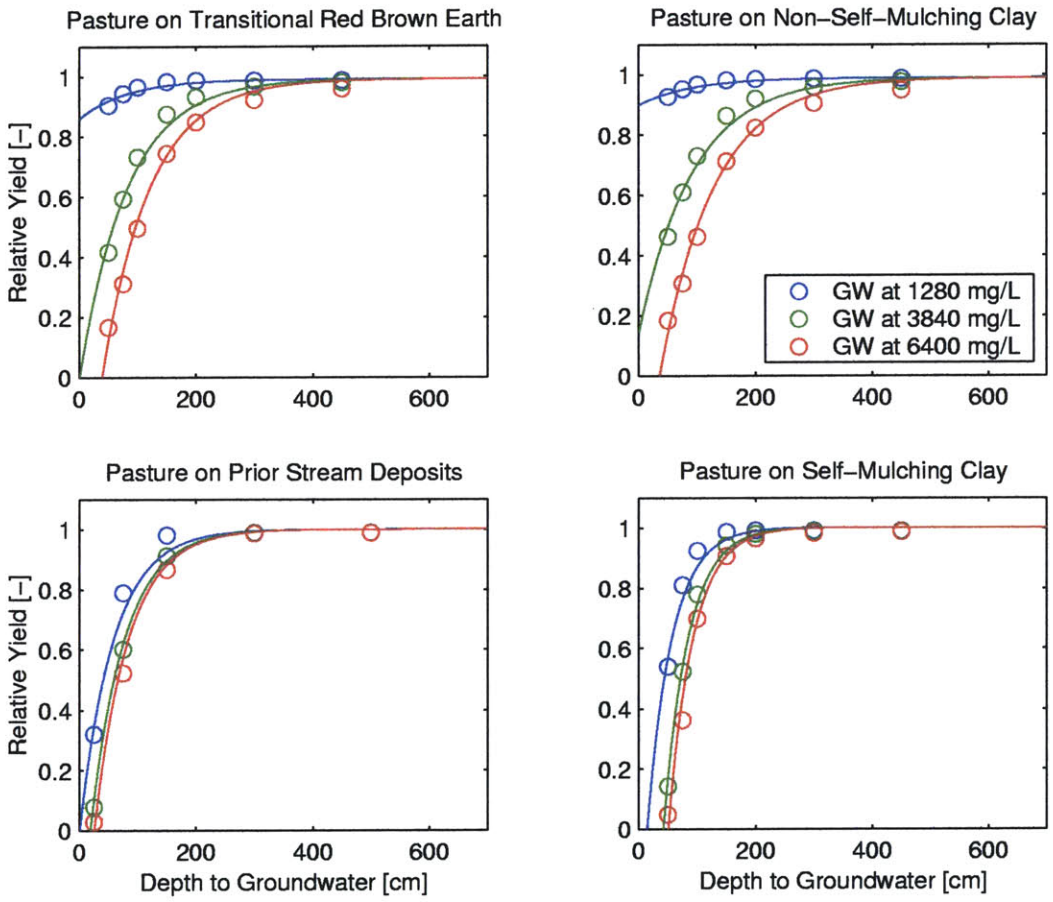

Figure A.7: Pasture Relative Yield Transfer Function

A.2.4 Fallow

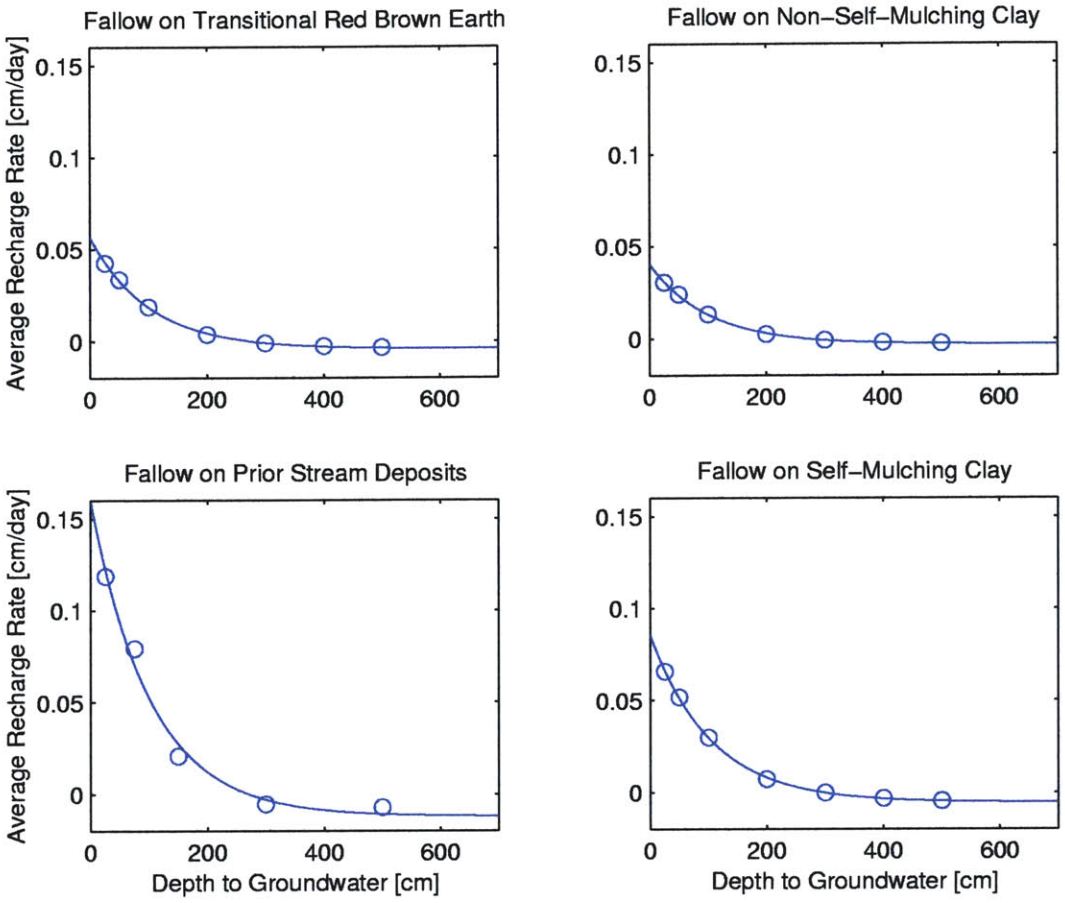

Figure A.8: Fallow Recharge Transfer Function

Bibliography

Ahmad, M. and Kutcher, G. P. [1992]. Irrigation planning with environmental considerations: A case study of Pakistan's Indus Basin, *Technical Paper No. 166*, World Bank, Washington, DC.

Allen, R. C. and Gisser, M. [1984]. Competition versus optimal control in groundwater pumping when demand is nonlinear, *Water Resources Research* **20**: 752–756.

Arad, A. and Evans, R. [1987]. The hydrogeology, hydrochemistry and environmental isotopes of the Campaspe River aquifer system, north-central Victoria, Australia, *Journal of Hydrology* **95**: 63–86.

Australian Bureau of Meterology [1988]. *Climatic Averages Australia*, Australian Government Publishing Service, Canberra.

Baumol, W. J. and Oates, W. E. [1988]. *The Theory of Environmental Policy*, 2nd edn, Cambridge University Press, Cambridge, UK.

Beecher, H. G. [1991]. Effect of saline water on rice yields and soil properties in the Murrumbidgee Valley, *Australian Journal of Experimental Agriculture* **31**: 819–23.

Bertsekas, D. P. [1999]. *Nonlinear Programming*, 2nd edn, Athena Scientific, Belmont, MA.

Bisschop, J., Chandler, W., Dudley, J. H. and O'Mara, G. T. [1982]. The Indus Basin model: A special application of two-level linear programming, *Math. Programming* **20**: 30–38.

Brooke, A., Kendrick, D., Meeraus, A. and Raman, R. [1997]. *Gams 2.25 Language Guide*, GAMS Development Corporation.

Brooks, R. H. and Corey, A. T. [1966]. Properties of porous media affecting fluid flow, *J. Irrig. Drainage Div., ASCE Proc.* **72**(IR2): 61–88.

Brown, C. M. [1989]. Structural and stratigraphic framework of groundwater occurrence and surface discharge in the Murray Basin, southeastern Australia, *BMR Journal of Australian Geology and Geophysics* **11**: 127–146.

California Department of Water Resources [1982]. The hydrologic and economic model of the San Joaquin Valley, *Bull. 214*, Calif. Dept. of Water Resources, Sacramento, CA.

Cardon, G. E. and Letey, J. [1992]. Soil-based irrigation and salinity management model: I Plant water uptake calculations, *Soil Sci. Soc. Am. J.* **56**: 1881–1887.

Chiang, R. Y. and Safonov, M. G. [1998]. *Matlab Robust Control Toolbox User's Guide, Version*, The Mathworks Inc.

de Marsily, G. [1986]. *Quantitative hydrogeology*, Academic Press.

Doorenbos, J. and Pruitt, W. O. [1977]. Guidelines for predicting crop water requirements, *Irrigation and Drainage Paper 24 (Revised)*, Food and Agricultural Organization of the United Nations, Rome.

Dorfman, R. [1969]. An economic interpretation of optimal control theory, *American Economic Review* **59**: 817–831.

Drud, A. S. [1996]. *CONOPT: A System for Large Scale Nonlinear Optimization, Reference Manual for CONOPT Subroutine Library*, ARKI Consulting and Development A/S, Bagsvaerd, Denmark.

Drury, L. W., Calf, G. E. and Dharmanasiri, J. K. [1984]. Radiocarbon dating of groundwater in Tertiary sediments of the eastern Murray Basin, *Australian Journal of Soil Research* **22**: 379–387.

Evans, W. R. and Kellett, J. R. [1989]. The hydrogeology of the Murray Basin, southeastern Australia, *BMR Journal of Australian Geology and Geophysics* **11**: 147–166.

Feddes, R. A., Bresler, E. and Neuman, S. P. [1974]. Field test of a modified numerical model for water uptake by root systems, *Water Resour. Res.* **10**: 1199–1206.

Feddes, R. A., Kowalik, P. J. and Zaradny, H. [1978]. *Simulation of field water use and crop yield*, Pudoc, Wageningen, The Netherlands.

Feinerman, E. and Knapp, K. [1983]. Benefits from groundwater management: magnitude, sensitivity, and distribution, *American Journal of Agricultural Economics* **65**: 703–710.

Ghassemi, F., Jakeman, A. J. and Nix, H. A. [1996]. *Salinisation of Land and Water Resources: Human Causes, Extent, Management and Case Studies*, CAB International, Wallingford Oxen, UK.

Gisser, M. and Sanchez, D. A. [1980]. Competition versus optimal control in groundwater pumping, *Water Resources Research* **16**: 638–642.

Gleick, P. H. [1993]. *Water in Crisis: A guide to the world's fresh water resources*, Oxford University Press, New York.

Glover, K. [1984]. All optimal Hankel-norm approximations of linear multi-variable systems and their L^∞-error bounds, *Int. J. Contr.* **39**: 1115–1193.

Hachett, S. A., Horner, G. L. and Howitt, R. E. [1991]. A regional mathematical programming model to assess drainage control policies, *in* A. Dinar and D. Zilberman (eds), *The Economics and Management of Water and Drainage in Agriculture*, Kluwer Academic, Boston.

Hahn, R. and Stavins, R. [1992]. Economic incentives for environmental protection: Integrating theory and practice, *American Economic Review* **82**(2): 464–472.

Hall, N., Poulter, D. and Curtotti, R. [1993]. ABARE model of irrigation farming in the southern Murray Darling Basin, *Research Report 94.4*, Australian Bureau of Agricultural and Resource Economics.

Harbaugh, A. W. and McDonald, M. G. [1996]. User's documentation for MODFLOW-96, an update to the U.S. Geological Survey Modular Finite-Difference Ground-Water Flow Model, *Open-File Report 96-485*, U.S. Geological Survey.

Harrison, S. and Chapman, L. [1999]. Productivity differences across New South Wales rice farms: Links to resource quality, *Conference Paper 99.2*, Australian Bureau of Agricultural and Resource Economics.

Hillel, D. J. [1991]. *Out of the Earth*, Free Press, New York.

Howitt, R. E. [1995]. Positive mathematical programming, *American Journal of Agricultural Economics* **77**(2): 329–342.

Humphreys, L., van der Lely, A., Muirhead, W. and Hoey, D. [1994]. The development of on-farm restrictions to minimize recharge from rice in New South Wales, *Australian Journal of Soil and Water Conservation* **7**(2): 11–20.

Jolly, I. D., Dowling, T. I., Zhang, L., Williamson, D. R. and Walker, G. R. [1997]. Water and salt balances of the catchments of the Murray Darling Basin, *Technical Report 37/97*, CSIRO Land and Water.

Jones, J. W. and Richie, J. T. [1990]. Crop growth models, *in* G. J. Hoffman, T. A. Howell and K. H. Solomon (eds), *Management of Farm Irrigation Systems*, American Society of Agricultural Engineers, St. Joseph, MI, pp. 63–85.

Langford-Smith, T. and Rutherford, J. [1966]. *Water and Land: Two Case Studies in Irrigation*, Australian National University Press, Canberra.

Lawson, S. and Webb, E. [1998]. Review of groundwater use and groundwater level behavior in the Lower Murrumbidgee Valley, *Technical Report 98/05*, Department of Land and Water Conservation, Leeton, New South Wales, Australia.

Lee, D., Howitt, R. E. and Marino, M. A. [1993]. A stochastic model of surface water quality: An application to salinity in the Colorado River Basin, *Water Resources Research* **29**: 3917–3923.

Lee, D. J. and Howitt, R. E. [1996]. Modeling regional agricultural production and salinity control alternatives for water quality policy analysis, *Amer. J. Agr. Econ.* **78**(1): 41–53.

Lefkoff, L. J. and Gorelick, S. M. [1990]. Simulating physical processes and economic behavior in saline, irrigated agriculture: Model development, *Water Resources Research* **26**(7): 1359–1369.

Letey, J. [1994]. Is irrigated agriculture sustainable?, *Soil and Water Science: Key to Understanding Our Global Environment*, Soil Science Society of America, Madison, WI, pp. 23–37.

Letey, J. and Dinar, A. [1986]. Simulated crop water production functions for several crops when irrigated with saline waters, *Hilgardia* **54**: 1–32.

Letey, J., Knapp, K. and Solomon, K. [1990]. Crop-water production functions, *in* K. Tanji (ed.), *Agricultural Salinity Assessment and Management*, American Society of Civil Engineers.

Maas, E. V. and Hoffman, G. J. [1977]. Crop salt tolerance–current assessment, *Journal of the Irrigation and Drainage Division* pp. 115–127.

Marshall, T. J., Holmes, J. W. and Rose, C. W. [1996]. *Soil Physics*, third edn, Cambridge University Press.

Meyer, W. S., Prapthapar, S. A. and Barrs, H. D. [1990]. Water flux to and from shallow water-tables on two irrigated soils, *in* E. Humphreys, W. A. Muirhead and A. van der Lelij (eds), *Management of Soil Salinity in South East Australia*, Australian Society of Soil Science, Riverina Branch, pp. 79–87.

MIA LWMP Taskforce [1998]. *MIA & Districts Community Land and Water Management Plan.*

Millington, R. J. and Quirk, J. M. [1961]. Permeability of porous soils, *Trans. Faraday Soc.* **57**: 1200–1207.

Molz, F. J. [1981]. Models of water transport in the soil-plant system: A review, *Water Resources Research* **17**: 1245–1260.

Moore, B. C. [1981]. Principal component analysis in linear systems: controllability, observability, and model reduction, *IEEE Trans. on Automat. Contr.* **AC-26**: 17–31.

Muirhead, W. A., Humphreys, E., Jayawardane, N. S. and Moll, J. L. [1996]. Shallow subsurface drainage in an irrigated vertisol with a perched water table, *Agricultural Water Management* **30**: 261–282.

Muirhead, W. A., Jayawardane, N. S., Humphreys, E. and van der Lelij, A. [1990]. Deep percolation under rice: Its importance, measurement and on-farm management, *in* E. Humphreys, W. A. Muirhead and A. van der Lelij (eds), *Management of Soil Salinity in South East Australia*, Australian Society of Soil Science, Riverina Branch, pp. 321–331.

National Research Council [1989]. *Irrigation-Induced Water Quality Problems*, National Academy Press, Washington, DC.

Negri, D. H. [1989]. The common property aquifer as a differential game, *Water Resources Research* **25**: 9–15.

Neuman, S. P., Feddes, R. A. and Bresler, E. [1974]. Finite element simulation of flow in saturated-unsaturated soils considering water uptake by plants, *Third Annual Report, Project No. A10-SWC-77*, Hydraulic Engineering Lab., Technion, Haifa, Israel.

Nieswiodomy, M. [1985]. The demand for irrigation water in the high plains of Texas, 1957-80, *American Journal of Agricultural Economics* **67**: 619–626.

NSW Department of Land and Water Conservation [1998a]. Murrumbidgee irrigation area & districts land & water management plan: Economic assessment of the upstream plan strategy, final report, *Technical Report CNR98.014*, Resource Economics and Sociology Unit, Parramatta, NSW.

NSW Department of Land and Water Conservation [1998b]. Economic assessment of MIA&D Integrated Land & Water Management Plan Strategy, downstream & integrated plan technical report, *Technical Report CNR98.013*, Resource Economics and Sociology Unit, Parramatta, NSW.

NSW Department of Land and Water Conservation [1998c]. Arcview datafiles of Murrumbidgee Valley rice fields 1993-1998. ESRI Corp. Arcview shapefile format.

NSW DLWC [1999]. Water trading development and monitoring.

Olsson, K. A. and Rose, C. W. [1978]. Hydraulic properties of a red-brown earth determined from in situ measurements, *Aust. J. Soil Res.* **16**: 169–180.

Postel, S. [1999]. *Pillar of Sand: Can the Irrigation Miracle Last*, W. W. Norton & Company, New York.

Punthakey, J. F., Somaratne, N. M., Prathapar, S. A., Merrick, N. P., Lawson, S. and Williams, R. M. [1994]. Regional groundwater modelling of the Lower Murrumbidgee River Basin: Model development and calibration, *Technical Report 94.069*, NSW Department of Water Resources, Parramatta, NSW.

Rhoades, J. D. [1997]. Sustainability of irrigation: An overview of salinity problems and control strategies, *Footprints of Humanity: Reflections on Fifty Years of Water Resource Developments*, CWRA, Lethbridge, Alberta, pp. 1–42.

Rugh, W. J. [1996]. *Linear System Theory*, Prentice Hall, Upper Saddle River, NJ.

Safonov, M. G. and Chiang, R. Y. [1989]. A Schur method for balanced model reduction, *IEEE Trans. on Automat. Contr.* **AC-34**: 729–733.

Shah, F., Zilberman, D. and Lichtenberg, E. [1995]. Optimal combination of pollution prevention and abatement policies: The case of agricultural drainage, *Environmental and Resource Economics* **5**: 29–49.

Sides, R. D., Prathapar, S. A. and Meyer, W. S. [1993]. Estimation recharge and monitoring the development of a water table mound from ponded rice, *Technical Memorandum 93/16*, Division of Water Resources, CSIRO.

Stannard, M. E. [1966]. North-east Riverine Plain physiographic units. Water Conservation and Irrigation Commission Map.

Taylor, C. R. and Howitt, R. E. [1993]. Aggregate evaluation concepts and models, *in* C. Carlson, D. Zilberman and J. Miranowski (eds), *Agricultural and Environmental Resource Economics*, Oxford University Press, Ontario.

van der Lelij, A. [1990]. Salt and water movement modeling, *in* E. Humphreys, W. A. Muirhead and A. van der Lelij (eds), *Management of Soil Salinity in South East Australia*, Australian Society of Soil Science, Riverina Branch, pp. 49–59.

van der Lely, A. [1998]. *Groundwater Conditions in the MIA and Districts, Annual Report 1998*, Murrumbidgee Irrigation.

van Genuchten, M. T. [1980]. A closed-form equation for predicting the hydraulic conductivity of unsaturated soil, *Soil Sci. Soc. Am. J.* **44**: 892–898.

van Genuchten, M. T. and Gupta, S. K. [1993]. A reassessment of the crop salt tolerance response function, *J. Indian Soc. Soil Sci.* **41**: 730–737.

van Genuchten, M. T. and Hoffman, G. J. [1984]. Analysis of crop salt tolerance data, *in* I. Shainberg and J. Shalhevet (eds), *Soil Salinity under Irrigation*, Springer-Verlag, Berlin, pp. 258–271.

van Schilfgaarde, J. [1996]. Irrigated agriculture: Is it sustainable?, *in* K. K. Tanji (ed.), *Agricultural Salinity Assessment and Management*, American Society of Civil Engineers, New York, chapter 28, pp. 584–594.

Vogel, T. and Cislerova, M. [1988]. On the reliability of unsaturated hydraulic conductivity calculated from the moisture retention curve, *Transport in Porous Media* **3**: 1–15.

Vogel, T., Huang, K., Zhang, R. and van Genuchten, M. T. [1996]. The HYDRUS code for simulating one-dimensional water flow, solute transport, and heat movement in variably-saturated media, Version 5.0, *Research Report No. 40*, U.S. Salinity Laboratory, Riverside, California.

Weitzman, M. L. [1974]. Prices vs. quantities, *Review of Economic Studies* **41**: 477–491.

Wichelns, D. [1999]. An economic model of waterlogging and salinization in arid regions, *Ecological Economics* **30**: 475–491.

Woolley, D. R. and Williams, R. M. [1978]. Tertiary stratigraphy and hydrogeology of the eastern part of the Murray Basin, NSW, *in* R. R. Storrier and I. D. Kelly (eds), *Proceedings of a symposium on the hydrogeology of the Riverine Plain of south-eastern Australia*, Australian Society of Soil Science, Riverina Branch, pp. 45–65.

World Bank [1992]. *Development and the Environment, World Development Report 1992*, World Bank, Washington, DC.

Worthington, V. E., Burt, O. R. and Brustkern, R. L. [1985]. Optimal management of a confined groundwater system, *Journal of Environmental Economics and Management* **12**: 229–245.

Wu, Q., Christen, E. and Enever, D. [1998]. Basinman: A water balance model for farms with subsurface pipe drainage and on-farm evaporation basin, *Technical Report 1/99*, CSIRO Land and Water, Griffith, NSW.