Multi-agent real-time decision making in water resources systems

by

Reetik Kumar Sahu

B.Tech., Mechanical Engineering
Indian Institute of Technology, Madras, 2013

S.M., Center for Computational Engineering
Massachusetts Institute of Technology, 2016

Submitted to the Center for Computational Engineering &
Department of Civil and Environmental Engineering

in partial fulfillment of the requirements for the degree of

Doctor of Philosophy in Computational Science and Engineering

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

September 2018

©2018 Massachusetts Institute of Technology. All rights reserved.

Signature of Author: Signature redacted
Reetik Kumar Sahu
Center for Computational Engineering
Department of Civil and Environmental Engineering
August 15, 2018

Certified by: Signature redacted
Dennis B. McLaughlin
H.M. King Bhumibol Professor of Water Resources Management
Thesis Supervisor

Accepted by: Signature redacted
Heidi Nepf
Donald and Martha Harleman Professor of Civil and Environmental Engineering
Chair, Graduate Program Committee

Accepted by: Signature redacted
Youssef M. Marzouk
Associate Professor, Aeronautics and Astronautics
Co-director, Center for Computational Engineering

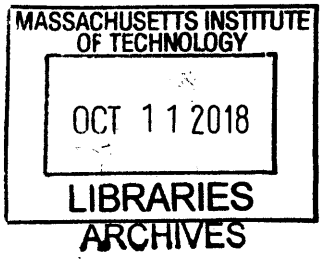

Multi-agent real-time decision making in water resources systems

by

Reetik Kumar Sahu

Submitted to the Center for Computational Engineering
& Department of Civil and Environmental Engineering
on August 15, 2018, in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy in Computational Science & Engineering

Abstract

Optimal utilization of natural resources such as water, wind and land over extended periods of time requires a carefully designed framework coupling decision making and a mathematical abstraction of the physical system. On one hand, the choice of the decision-strategy can set limits/bounds on the maximum benefit that can be extracted from the physical system. On the other hand the mathematical formulation of the physical system determines the limitations of such strategies when applied to real physical systems. The nuances of decision making and abstraction of the physical system are illustrated with two classical water resource problems: optimal hydropower reservoir operation and competition for a common pool groundwater source. Reservoir operation is modeled as a single agent stochastic optimal control problem where the operator (agent) negotiates a firm power contract before operations begin and adjusts the reservoir release during operations. A probabilistic analysis shows that predictive decision strategies such as stochastic dynamic programming and model predictive control give better performance than standard deterministic operating rules. Groundwater competition is modeled as a multi-agent dynamic game where each farmer (agent) aims to maximize his/her personal benefit. The game analysis shows that uncooperative competition for the resource reduces economic efficiency somewhat with respect to the cooperative socially optimum behavior. However, the efficiency reduction is relatively small compared to what might be expected from incorrect assumptions about uncertain factors such as future energy and crop prices. Spatially lumped and distributed models of the groundwater system give similar pictures of the inefficiencies that result from uncooperative behavior. The spatially distributed model also reveals the important roles of the geometry and density of the pumping well network. Overall, the game analysis provides useful insight about the factors that make cooperative groundwater management beneficial in particular situations.

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

Acknowledgments

I owe my deepest gratitude to Prof. Dennis McLaughlin for his unending support and also for being a constant source of motivation throughout my time here at MIT. I am especially thankful for his ever-approachable nature and hope to carry on his drive to bring out something new and important through our research.

I would also like to thank my committee members, Prof. Saurabh Amin, Prof. Kenneth Strzepek and Prof. Charles Harvey for their invaluable feedback and mentoring to become a better researcher.

Special thanks to my research group members Anjuli Jain Figueroa and Tiziana Smith and to the Parsons community for the dynamic conversations making the my stay at Parsons a priceless memory.

I am also grateful to my several friends who have helped me find a home here at MIT- Saurabh Gandhi, Ankit Shah, Rajesh Sridhar, Sami Khan, Chaitanya Talnikar, and other close friends in MIT Indian graduate community for all our discussions that has been a constant source of happiness, intrigue and fun.

Special mention and thanks to the MIT/ Boston Metro Badminton group: Jeffrey Mei, Daphne Chang, Neil Davis and others members of the group who have been a major part of my time outside MIT.

I would also like to thank Kate Nelson (CCE), Jim Long (CEE) and Ed Carlevale (CEE) for all their help with the administrative work and making everyday at MIT a smooth experience.

Last, but not the least, I am highly indebted to my parents, Pradip and Meena, my brother Pratik and sister-in-law Sweta for their continual encouragement and love. I am also thankful to my fiancé Uma Maheswari Selvaraj for her constant support throughout this PhD experience.

Finally thanks to each and everyone who were directly or indirectly helped me during the course of the project.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

1.1 Motivation

Large-scale structures and systems are built around natural resources like water, wind, land, fossil fuels and several others to extract benefits for human use. Operating around such physical systems requires a careful understanding of its underlying dynamics and the effects of our actions on such systems. However these sytems are stresses due to increasing demand and standard of living of the growing human population. Problems arising from these shortages are further aggravated due to increasing seasonal variability in weather patterns and disruptions in the natural cycles of resources from natural and anthropogenic activities.

Therefore, rather than maximizing the extraction of economic benefit from these natural resources, the focus of research has shifted towards developing sustainable and efficient use of these natural resources. As a decision-maker one has to balance his/her inclination to extract greater benefits at the current time with the possibility of incurring lower benefits or irreversible losses in the future.

The primary focus of this research is to implement numerical solutions of dynamic optimization techniques to permit the use of a more realistic physical system. On one side, we analyze the and compare the performance and implications of different decision making strategies to a given physical model. And on the other side, we compare effect of model description on its derived decision behavior. Such an analysis requires first a mathematical formulation of the physical process and the limits in the natural system that act as constraints. The second part focuses on building optimization models that can perform on such complex systems. These optimization models coupled with optimal control techniques like dynamic programming and game theory, can help us better understand and quantify resource utilization.

The implementation of the above techniques is illustrated with two water resource utilization problems: optimal hydropower operation and competition for a groundwater resource. The reservoir operation is a modeled as a single agent optimal control problem where it must negotiate an appropriate energy contract that it can reliably provide and operate the reservoir in real time adapting to uncertainty in the reservoir inflow and other environmental constraints. The competition for a groundwater

resource is modeled as multi-agent optimization problem where each agent aims to maximize his/ her personal objective. A dynamic game approach is used to analyze the behavior if such a system. Applying a more realistic description of the physical system yields results that are qualitatively different from previous solutions obtained from the simpler abstractions of the physical system.

1.2 Single agent hydropower operation

Revenues from hydropower generation often depend on the operatorâĂŹs ability to provide firm power in the presence of uncertain inflows. The primary options available for optimizing revenue are negotiation of a firm power contract before operations begin and adjustment of the reservoir release during operations. Optimization of the contract and release strategy is closely coupled and most appropriately analyzed using stochastic real-time control. Here we use an ensemble-based approach to stochastic optimization that provides a convenient way to construct non-parametric revenue probability distributions to explore the implications of inflow uncertainty. The ensemble approach makes it possible to readily compare revenue distributions for different contract selection-reservoir operations strategies. These distributions and related spill performance statistics reveal the distinctive risk characteristics of each strategy. They suggest that predictive operating strategies such as stochastic dynamic programming and model predictive control can give significantly better performance than standard deterministic operating rules. The performance obtained from batch optimization with perfect inflow information establishes a convenient upper bound on potential revenue and provides a baseline for assessing the significance of differences between real-time operating strategies. Sensitivity analysis indicates that the benefits of predictive operational strategies are greatest for reservoirs with medium non-dimensional residence times and less important for reservoirs with large residence times. These strategies can better accommodate spill penalties than standard operating rules. Overall, probabilistic analysis of the coupled hydropower contract-operations problem provides a realistic way to assess revenue and risk for reservoirs that must provide firm power when inflows are uncertain.

1.3 Multi-agent groundwater common pool management: A lumped model approach

The groundwater common pool resource problem is examined with special attention given to the hydrologic realism of the problem formulation. Groundwater depletion is considered for an uncooperative open-loop Nash equilibrium, a sub-game perfect feedback Nash equilibrium, and a myopic pumping strategy as well as a cooperative social optimum strategy that establishes a baseline for comparing the economic efficiencies of the alternatives. System dynamics are described with a two-compartment model that includes both groundwater recharge and flow to or from a river. This model provides realistic steady-state solutions for both pumped and unpumped conditions

and meets the requirements needed to obtain dynamic closed-form discrete-time solutions for all of the selected management strategies. The availability of these solutions makes it possible to quantify the magnitude of the strategic externality that results from uncooperative aquifer exploitation as well as the effects of player asymmetry, the number of players, and hydrologic variability. The formulation can also quantify the effect of stock limitations for the social optimum and myopic cases. These factors are all considered in a simple example that captures several key hydrologic features of interest in practical groundwater depletion problems. Generally speaking, the inefficiency and strategic externality due to uncooperative depletion are small relative to the total revenue, with economic efficiency decreasing progressively from the social optimum to the myopic strategy. The results show how hydrologic considerations previously neglected in game theoretic analyses can influence player behavior and the state of the groundwater resource.

1.4 Multi-agent groundwater common pool management: A spatially distributed approach

The common pool resource problem is reexamined with a spatially distributed groundwater model. System dynamics using such a model can capture the inter-spatial interactions between the players' pumping wells and the boundaries. With suitable assumptions the groundwater depletion can be reconsidered for a socially optimum, sub-game perfect/ feedback Nash equilibrium and a myopic pumping strategy. The economic efficiencies are measured with the socially optimum as the baseline strategy. The strategic externality from non-cooperative behavior is measured for different pump well configurations and player asymmetry. The results show flow behavior and flux distribution neglected in a lumped model are important to consider when considering water budgets and policies.

Chapter 2

An ensemble optimization framework for coupled design of hydropower contracts and real-time reservoir operating rules

2.1 Introduction

Hydroelectricity contributes 71% of global renewable electrical energy and 16% of total global energy demand. Much of this energy is sold to institutional and industrial clients under contractual agreements. A typical contract specifies a price to be paid per unit energy for firm (continuous) power output or for a firm amount of energy to be delivered over a designated time period. Such firm power agreements are complicated by the fact that hydropower is an intrinsically uncertain energy source that depends on variable river inflows. When determining a contract power or energy target, the operator must trade off the risk of having to purchase make-up power during low-inflow periods vs. the risk of forgoing potential income during high inflow periods. Both of these situations can reduce net present value revenue if the contract target is not properly chosen. Hydropower revenue depends not only on the contract terms, but also on the inflow characteristics, facility specifications (e.g. reservoir geometry and turbine capacity), electricity demand and pricing, and the operating rule used to determine the reservoir release at any given time. This paper emphasizes connections between inflow uncertainty and real-time operations, recognizing that management of risk in hydropower generation is a much broader topic that involves many other sources of uncertainty [16], [73], [23].

When developing a strategy for managing inflow uncertainty, it is best to determine the contract and the operating rule together. The optimum contract depends in part on the operational strategy used to deal with inflow variability and the optimum operating rule depends in part on the requirements specified in the contract. This chapter addresses the coupled contract-operational design problem by simultaneously deriving an optimum firm energy value and an associated real-time operating rule.

Together, these maximize expected revenue over the contract period for given inflow statistics, contract prices, and other inputs.

Optimization techniques have been used extensively by the reservoir operations community, most often to derive operating rules than energy contracts. [82], [79], [41], and [61] provide comprehensive literature reviews on methods to optimize reservoir operation. In practice, most hydropower reservoirs are managed with deterministic operating rules that fall under the umbrella of a Standard Operating Policy (SOP) with hedging ([52], [75], [83]). These rules typically relate the current reservoir release to the current reservoir storage and do not attempt to predict or adjust for future inflow variations.

Reservoir operators may be able to extract more benefit than can be achieved with Standard Operating Policies if they adopt predictive rules that rely on probabilistic models of future inflows. The gold standard of this approach is Stochastic Dynamic Programming (SDP) [11]. The SDP method has been applied to the control of a single reservoir and to networks of multiple reservoirs ([72],[84],[15],[17],[32],[33], [80],[82]). The popularity of SDP lies in its flexibility to accept a variety of objectives, constraints (equality and/or inequality), and random disturbance models. Its main limitation is its computational complexity, which grows very quickly with the number of state and control variables used to describe the reservoir system (the so-called *curse of dimensionality*). This reflects the fact that SDP derives a general control law that optimizes the complete trajectory of future releases for any feasible current state. Several approximate SDP techniques have been developed to deal with the method's computational demands. These take advantage of distinctive structural features applicable to reservoir operations problems ([4],[12],[11],[42],[15]).

Stochastic Model Predictive Control (SMPC) ([26],[47],[62]) method is a limited look ahead real-time optimization technique that also relies on probabilistic inflow models. SMPC is generally less computationally intensive than SDP and is able to readily handle complex constraints. This reflects the fact that SMPC optimizes the current release only for a particular (observed) current state. SMPC plays an important role in process control where efficiency requires operating the system near specified bounds on the state and the control. Examples of relevant SMPC applications include process control [27], reservoir operations ([5],[43]), irrigation ([53],[76]), and supply chain management ([55],[60]). Here we consider the performance obtained with all three of the reservoir operating rules mentioned above (SOP, SDP, and SMPC) when they are coupled with a contract optimization procedure.

2.2 Formulation of the reservoir problem

In this paper we illustrate the concept of coupled contract-operational design by adopting a simple contract structure that serves to illustrate the connection between a firm power commitment and the design of predictive real-time operating rules. We suppose that an operator of a reservoir designed primarily for hydropower production sells power to a single consumer, according to a contract agreed upon before operations start. The consumer desires a dependable source of firm power, defined here as a

specified power value generated continuously throughout a multiyear contract period. An example is an industrial client, such as an aluminum manufacturer, who has a continuous demand for fixed amount of power. The customer agrees to pay a specified unit price ($\$MWhr^{-1}$) for the firm power, negotiated at the start of the contract period and held fixed until the end of the period. Recognizing that it may not always be possible to meet a particular firm power value when reservoir inflows are variable, the contract stipulates that shortfalls be covered by the consumer, who buys makeup power at the market price and passes on a fixed unit charge to the operator. If the market price is below the shortfall charge the consumer benefits. If it is above then the operator benefits. Similarly, the consumer agrees to buy surplus power for a negotiated fixed price. If this price is above the market price the operator benefits. If it is lower then the consumer benefits. The net result depends on the relationship between variations in market prices and variations in inflows at the reservoir. Our contract arrangement has the advantage of providing the reservoir operator with a predictable pricing structure so the major source of operational uncertainty is inflow variability rather than price variability. We quantify this uncertainty with revenue probability distributions that apply for several different operating strategies.

Other types of contracts may able to better manage risk and perhaps improve the operator's expected revenue. Examples relevant to hydropower applications are discussed in the financial risk literature ([23],[16],[48]). These include methods that provide various forms of insurance to protect the operator from uncertainty. The advantage of our contract is that it insulates the operator from price fluctuations and focuses attention on real-time operations. The ensemble-based optimization approach introduced here is intended to provide a general framework that can be used to examine the implications of different operational strategies as well as different contractual arrangements. Our fixed price contract provides a convenient way to show how this framework can be applied.

In a discrete time problem formulation the firm power value can be interpreted as an equivalent firm energy generated over a constant time step. Our coupled contract-operational design optimization focuses on two types of decision variables: 1) the firm energy value negotiated with the consumer and 2) the reservoir releases at a set of regularly spaced decision times throughout the contract period. The releases may be determined from an operating rule that is derived as part of the optimization process. The optimum firm energy value and operating rule maximize the operator's expected present value revenue in the presence of uncertain inflows, subject to relevant physical constraints.

This coupled stochastic optimization problem can be solved with an iterative algorithm that starts with an initial value for the firm energy value and an initial operating rule based on this value. The algorithm evaluates the resulting revenue, adjusts the contract energy to increase revenue, derives a new operating rule, and again evaluates the revenue, continuing until the process converges (Figure 2-1).

It is helpful to describe the contract in a mathematical form suitable for the optimization. Suppose for a given firm contract energy E_c that the reservoir generates actual energy E_k over the unit time interval $[t_k, t_(k+1)]$. The revenue obtained over this interval is determined by the piecewise linear concave revenue function illustrated

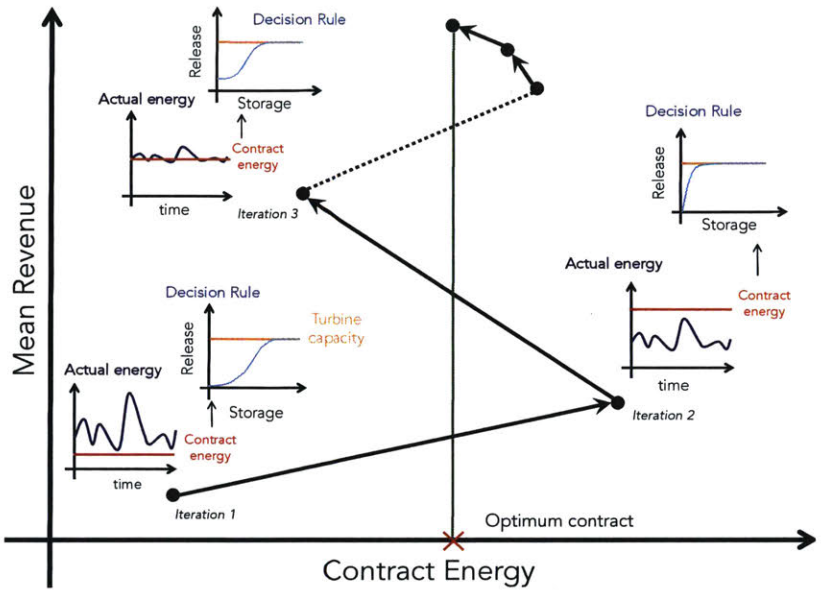

Figure 2-1: Iterative search for optimum energy contract. The energy contract E_c proposed at each iteration requires a new decision rule (turbine release vs. storage) to maximize revenue.

in Figure 2-2:

$$g(E_k, E_c) = \alpha_1(E_k - E_c) + \alpha_c E_c \quad if\ E_k E_c \tag{2.1}$$

$$g(E_k, E_c) = \alpha_2(E_k - E_c) + \alpha_c E_c \quad if\ E_k > E_c \tag{2.2}$$

Where E_c is the firm contract energy to be generated in each time interval during the contract period and $\alpha_1 > \alpha_c > \alpha_2$ are coefficients that define the price per unit energy in ($\$MWhr^{-1}$) that applies for different situations. The term $\alpha_c E_c$ is the revenue obtained if the contract is exactly satisfied i.e. $E_k = E_c$. If E_k is greater than E_c, there is a surplus and the operator sells the additional energy $E_k - E_c$ at a lower rate $\alpha_2 < \alpha_c$. If E_k is less than E_c, there is a shortfall and the operator must purchase makeup energy $E_c - E_k$ at a higher rate $\alpha_1 > \alpha_c$. As mentioned earlier, we assume that the prices α_c, α_1 and α_2 are fixed, but that E_c is a decision variable. In contracts based on spot rather than fixed surplus and makeup power prices, variability in α_1, α_c,and α_2 could be an important contributor to revenue uncertainty. This extension can be incorporated in the ensemble approach outlined here once the alternative contractual arrangement is precisely defined.

We suppose that the contract period extends from times t_1 to t_K and is divided into $K-1$ intervals of fixed duration Δt. The firm energy E_c needs to be known at the beginning of the contract period when the contract is negotiated. By contrast, the reservoir release u_k is most appropriately determined in real-time over each discrete time interval in the contract period, as illustrated in Figure 2-3. Real-time operation is important since it takes into account unanticipated variations in inflow and storage.

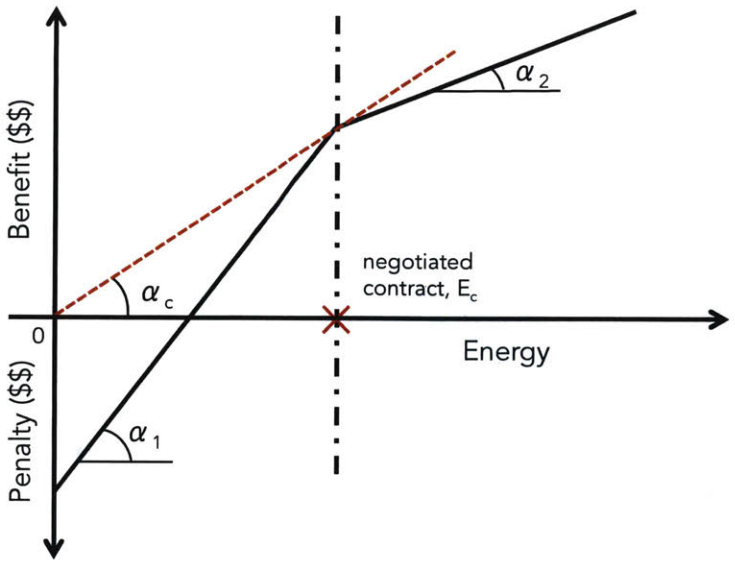

Figure 2-2: Piecewise linear concave revenue function. The slope of the red-dotted and black lines shows how the unit revenue at the contract energy compares to the unit cost of makeup power at lower energy values and the unit revenue of surplus energy at higher energy values

To examine the real-time aspect in more detail we need to characterize the dynamic behavior of the reservoir system, which is described by the system states, releases, and energy output. These variables can be related by a set of stochastic constraints. For the single reservoir hydropower problem considered here the state vector x_k is partitioned into a scalar reservoir storage S_k, observed at time t_k, and a vector of states ψ_k that collectively describe an inflow time series model that could be estimated from observed inflow data using system identification techniques [44]. The associated state equations are:

$$x_{k+1} = f(x_k, u_k, \omega_k) \tag{2.3}$$
$$x_k = [S_k, \psi_k]^T \tag{2.4}$$
$$S_{k+1} = f_S(S_k, \psi_k, u_k, \omega_k) \tag{2.5}$$
$$\psi_{k+1} = f_\psi(\psi_k, \omega_k) \tag{2.6}$$
$$I_k = M(\psi_k) \tag{2.7}$$

Here I_k is the total reservoir inflow over the time interval $[t_k, t_{k+1}]$, observed at t_k. This inflow is related to the time series model state ψ_k through a specified function $M(.)$. The scalar u_k (the control variable) is the total reservoir release over $[t_k, t_{k+1}]$ (specified by the operator at t_k) and ω_k is a sequence of independent random disturbances that drives the time series model. The time series model is used to predict inflows for the predictive operating rules considered in our example. Specific

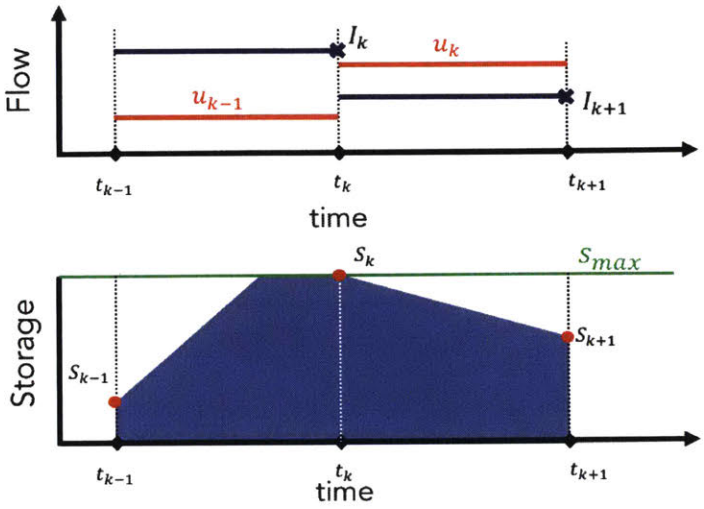

Figure 2-3: Example representation of discrete reservoir variables defined over two consecutive time intervals. The bottom panel shows the piecewise linear storage state (S) over each interval. The top panel shows the piecewise constant turbine release (u) and inflow (I) over each interval, with the inflow measurement (I) observed at the end of the interval.

options for this model and its associated functions and variables are discussed in Section 2.4.

The storage state equation is a mass balance expression that neglects evaporation and seepage but includes spills:

$$S_{k+1} = f_S(S_k, \psi_k, u_k, \omega_k) = S_k + \Delta t[I_{k+1}(\psi_k, \omega_k) - u_k] - Z_k \qquad (2.8)$$
$$@k = 0; S(0) = S_0$$

Where the expressions in eqns 2.6 and 2.7 can be used to write I_{k+1} in terms of the state vector ψ_k and disturbance ω_k. This state equation is used by the predictive operating rules to forecast storage from a particular predicted inflow sequence. The reservoir spill Z_k over $[t_k, t_{k+1}]$ is given by an additional constraint:

$$Z_k = \max\{S_k + \Delta t[I_{k+1} - u_k] - S_{max}, 0\} \qquad (2.9)$$

Where $S_{max}(L^3)$ is the reservoir capacity. The energy E_k generated by releases over $[t_k, t_{k+1}]$ is:

$$\begin{aligned} E_k &= \phi(u_k, h_k, h_{k+1})) = u_k \int_{t=k}^{k+1} H(f_S)dt \qquad (2.10) \\ E_{max} &= \phi(\bar{I}, h_{max}, h_{max}) \end{aligned}$$

The reservoir head h_k at time t_k is related to the storage by a specified function $H(.)$ that depends on the reservoir geometry:

$$h_k = H(S_k); \quad h_{max} = H(S_{max}) \tag{2.11}$$

The controlled release is constrained to be no greater than the turbine capacity u_{max}

$$u_k \leq u_{max} \tag{2.12}$$

For purposes of this study, the reservoir capacity, the head-storage function, and the turbine capacity are all assumed to be given. Note that fluxes are defined over K time intervals indexed by $k = 0 : K-1$ and states are defined at $K+1$ discrete times indexed by $k = 0 : K-1$. The time series notation can be made more concise if the entire sequence of releases defined through any time h_k is represented by the vector $u_(0, 1, , k-1) = u_(0 : k-1)$. Similar notation is used for sequences of other variables. The desired solution to the real-time operations problem maximizes the following expected present value objective, which measures performance over the contract period $[t_0, t_K]$ for a given sequence of releases $u_(0 : K-1)$, a given initial state x_0 , and a given firm energy value E_c:

$$J(u_{0:K-1}, x_0, E_c) = \mathrm{E}_{\omega_{0:K-1}} \left\{ \sum_{k=0}^{K-1} (1+r)^{-k} [g(E_k(u_k, x_k, x_{k+1}), E_c) - \alpha_Z Z_k(u_k, x_k)] + g_K(x_K) \right\} \tag{2.13}$$

Dependence on the vector of random inflow disturbances $\omega_{0:K-1}$ is removed by the expectation operation $\mathrm{E}_{\omega_{0:K-1}}$. The first term in the objective function expression is the present values of the hydropower revenue. The second term $\alpha_Z Z_k$ penalizes reservoir spills that can cause downstream flooding. The final term $g_k(x_K)$ (the salvage value) assigns a prescribed benefit to reservoir storage at the final time. This prevents the control strategy from emptying the reservoir at the end of the contract period. Specification of the spill and salvage value terms is discussed in more detail in Section 2.4.

The objective given in eqn 2.13 could be maximized simultaneously with respect to the variables $u_{0:K-1}$ and E_c, using the methods of mathematical programming, imposing the constraints identified above. Since the contract must be determined before operations begin, at t_0, a simultaneous optimization of $u_{0:K-1}$ and E_c would require the entire release history to also be derived at t_0, before there are any observations of the actual states (open-loop control). Better revenue can generally be obtained if the contract is determined at the initial time but the releases are determined in real-time, as observations of the states become available (closed-loop control). This is possible if the release at each time is derived directly from the observed state, as specified by a closed loop operating rule (or decision function) of the following form:

$$u_k = \mu_k(x_k) \quad k = 0 : K-1 \tag{2.14}$$

If eqn 2.14 is substituted into eqn 2.13 the objective J can be written as a func-

tional $J_{\mu_{0:K-1}}(x_0, E_c)$ that maps the K decision functions $\mu_{0:K-1}$ to the scalar revenue.

$$J_{\mu_{0:K-1}}(x_0, E_c) = \mathrm{E}_{\omega_{0:K-1}}\{\sum_{k=0}^{K-1}(1+r)^{-k}[g(E_k(x_0, \mu_{0:k}, \omega_{0:k}), E_c) - \alpha_Z Z_k(x_0, \mu_{0:k}, \omega_{0:k})] + g_K(x_0, \mu_{0:K-1}, \omega_{0:K-1})\} \tag{2.15}$$

This is the real-time optimal control form of the optimization objective given in eqn 2.13.

Note that the revenue, spill, and salvage terms are all random by virtue of their dependence on the random disturbance vector $\omega_{0:K-1}$. In our ensemble implementation of the stochastic optimal control problem, many random samples (or replicates) of this vector are drawn from a population determined by the statistics of the inflow time series model. Each inflow replicate gives a corresponding sample for each of the three terms in the objective and for the objective as a whole. The objective replicates provide equally likely predictions of the system performance for a given firm energy and decision strategy. The expected objective value is estimated by the arithmetic average of these replicates.

A real-time formulation of the operational part of the coupled optimization problem makes it possible to more precisely describe the iterative procedure outlined in figure 2-1. If the current iterates (for iterations $l = 1, , L$) of the decision function and firm energy value are $\mu^l_{0:K-1}$ and E^l_c the new decision strategy $\mu^{l+1}_{0:K-1}$ is obtained by maximizing with respect to all decision functions that satisfy constraints eqns 2.4-2.12.

$$\mu^{l+1}_{0:K-1} = \arg\max_{E_c} J_{\mu_{0:K-1}}(x_0, E^l_c) \tag{2.16}$$

This real-time optimal control sub-problem can be solved with the SDP, SMPC, and PI methods described in more detail in Section 2.3. Then the new firm energy value is obtained by maximizing $J_{\mu^{l+1}_{0:K-1}}(x_0, E_c)$ with respect to the scalar E_c:

$$E^{l+1}_c = \arg\max_{E_c} J_{\mu^{l+1}_{0:K-1}}(x_0, E_c) \tag{2.17}$$

This scalar optimization sub-problem can be readily solved with a one-dimensional search procedure (e.g. the Newton Rhapson method). The iteration can be initialized with a plausible firm energy value, such as the energy that could be generated with a constant inflow somewhat less than the observed mean.

We denote the converged decision function and firm power by $\mu^*_{0:K-1}$ and E^*_c . We are unaware of a convergence proof for this algorithm but it has always converged in less than 20 iterations in the many sensitivity analyses we have performed for all of the predictive operating rules considered in Section 2.4. The iterates are well-constrained by the inflows and by the physical limitations of the reservoir system and all discontinuities (e.g. the reservoir spill expression) are approximated by locally smooth functions. Our experience has been that these factors lead to quick and

reliable convergence.

The random inflow disturbance replicates generated in the iteration outlined above are used to guide the search procedure. In a practical application, the resulting optimum decision function and firm energy are used to determine the actual release from the reservoir. The corresponding actual inflow disturbance sequence will generally be different from any of the replicates used in the iteration. It is useful to quantify how well the reservoir system might work in such a situation. Since we do not know the actual inflows in advance such a performance assessment should account for uncertainty by considering a range of possible actual inflow disturbances. The framework for this assessment can be formulated in terms of the actual objective $J^a_{\mu^*_{0:K-1}}$, which depends on the actual inflow disturbance vector $\omega^a_{0:K-1}$ and the actual initial state as follows:

$$J^a_{\mu^*_{0:K-1}}(\omega^a_{0:K-1}, x^a_0, E^*_c) = \sum_{k=0}^{K-1}(1-r)^{-k}[g(E_k(x^a_0, \mu^*_{0:k}, \omega^a_{0:k}), E^*_c) - \alpha_Z Z_k(x^a_0, \mu^*_{0:k}, \omega^a_{0:k}] + g_k(x^a_0, \mu^*_{0:K-1}, \omega^a_{0:K-1}) \tag{2.18}$$

Here E_c and the decision functions $\mu^*_{0:K-1}$ from eqn 2.16 have been identified from the optimization procedure and can be considered given. At the initial time, before the inflows are observed, $\omega^a_{0:K-1}$ can be viewed as a random sequence sampled from the same population as the sequence that appears in eqn 2.15. If x^a_0 is also unknown at the initial time it can also be treated as a random variable with a specified distribution. We call a collection of $\omega_(0:K-1)^a$ and x^a_0 samples a *meta-ensemble* to distinguish it from the ensemble $\omega_{0:K-1}$ used in the iterative search procedure.

If eqn 2.18 is evaluated for a meta- ensemble of $\omega^a_{0:K-1}$ and x^a_0 samples we can derive the probability distribution of the actual present value revenue before inflows are actually observed. This distribution can be used to compute various revenue statistics such as the mean, upper quantile, etc. The process is carried out for selected decision rules in section 2.4.

2.3 Options for deriving real-time decision strategy

The options for deriving the operating rule $\mu_k(x_k)$ use different methods to relate the current release to the current state. This section reviews some of the most promising alternatives.

2.3.1 Stochastic Dynamic Programming

Stochastic dynamic programming (SDP) provides a comprehensive approach for deriving real-time operating rules before real-time operations begin, without simplifying assumptions. In the discrete time version used here this method divides the real-time control problem of 2.14 and 2.15 into a sequence of K nested sub-problems that

are solved with a recursion ([10],[11]). Each subproblem optimizes a time-dependent objective (the benefit-to-go) from a particular time to the end of the contract period. The objective for subproblem k, which is associated with time interval $[t_{k-1}, t_k]$ (commonly called Stage k) is:

$$J_{SDP,k}(x_k, E_c) = \max_{\mu_k(x_k)} [\mathrm{E}_{\omega_k}\{g[E_k(x_k.\mu_k(x_k), \omega_k), E_c] - \alpha_Z Z_k(x_k, \mu_k(x_k), \omega_k) + (1+r)^{-1} J_{SDP,k+1}(x_{k+1}, E_c)\}] \tag{2.19}$$

The problems are nested because sub-problem k depends on the solution of sub-problem $k+1$. The solution is computed with a backward recursion that moves stage by stage from the final to initial contract times. A decision function $\mu_k(x_k)$ is derived and stored for sub-problem k (for $k = K-1,,0$), for a given E_c. The recursion is initialized at $k = K$:

$$J_{SDP,K}(x_K, E_c) = g_K(x_K) \tag{2.20}$$

Note that the objective $J_{SDP,0}(x_0, E_c)$ obtained at the end of the recursion is equal to optimal revenue objective $J^*_{SDP,0}(x_0, E_c)$ defined in 2.17. Also, the state equation can be used to express the term $J_{SDP,k+1}(x_{k+1}, E_c)$ appearing in 2.19 as a functional that depends on $x_k, \mu_k(x_k), \omega_k$ and E_c. When the recursion is complete, the decision functions for all intervals are available and can be used to compute releases from actual observations in a forward real-time sweep (for $k = K-1,,0$).

The maximization over $\mu_k(x_k)$ of the expected revenue in eqn 2.19 gives the optimal release Stage k for any given value of the current state x_k. In practice, the state vector is usually discretized into a finite number of grid points and the optimum release value u^*_k is found at each of these points by maximizing the argument of eqn 2.19, with E_c fixed. The releases at the grid points are interpolated to give a decision function $\mu_k(x_k)$ that applies at any feasible value of the state ([17],[35]). The expectation operation appearing in eqns 2.18 and 2.19 is approximated by the mean over an ensemble of synthetically generated ω_k samples, as discussed in Section 2.2.

Some distinctive aspects of the dynamic programming approach include:

- the decision rules for all times are derived prior to the start of operations but each reservoir release is derived in real-time, after the current state is observed;
- the decision function in our formulation depends on the energy contract;
- the computational effort grows rapidly as the problem size increases. If N_{x_t}, N_{u_t} and N_{ω_t} are the number of discretized states, controls and inflow disturbances and the optimization horizon is K time steps, then the SDP algorithm requires $KN_{x_t}N_{u_t}N_{\omega_t}$ functional evaluations of the objective function;
- performance is dependent on the accuracy of the predictive inflow and storage models (the stochastic state equations)
- the algorithm implicitly accounts for the information provided by future measurements by relying on conditional probabilities that determine the likelihood

of a transition from a particular observed state at t_k to another state at t_{k+1}.

The computational demands of SDP tend to limit its application to problems with relatively small state vectors. In the hydropower operations context this implies that the problem needs to include only a few reservoirs and/or low dimensional inflow models.

2.3.2 Stochastic model predictive control

Stochastic model predictive control (SMPC) derives the optimal release u_k^* at each decision time by maximizing expected revenue over a limited duration window extending into the future. The complete series of reservoir releases is computed by carrying out a new optimization at every time step rather than using a pre-computed decision rule. The objective for Problem k originating at t_k is the present value revenue from t_k to t_K, based on eqn 2.15 and written directly in terms of releases rather than in terms of a decision function:

$$J_{SMPC,k}(u_{k:k+w-1}, x_k, E_c) = \mathrm{E}_{\omega_{k:k+w-1}}\{\sum_{i=k}^{k+w-1}(1+r)^{-i}[g[E_i(x_k, u_{k:i}, \omega_{k:i}, E_c] - \alpha_Z Z_i(x_i, u_{k:i}, \omega_{k:i}] + g_{k+w}(x_k, u_{k:k+w-1}, \omega_{k:k+w-1})\} \tag{2.21}$$

The expectation operator is approximated by the mean over an ensemble of synthetically generated samples $\omega_{k:k+w-1}$. The optimization is carried out over a moving window of length $wK - k$ time steps. This window spans the interval $[t_k, t_{k+w}]$.

An optimal release sequence over the current SMPC window is obtained by maximizing $J_{SMPC,k}(x_k, E_c)$ with respect to the releases:

$$u^*_{k:k+w-1} = \arg\max_{u_{k:k+w-1}} J_{SMPC,k}(u_{k:k+w-1}, x_k.E_c) \tag{2.22}$$

Although this optimization gives an entire sequence of optimal releases over the current time horizon, only the first release u_k is actually applied to the reservoir system (at t_k) since the remaining releases are recomputed at t_{k+1} when a new value of the state x_{k+1} is observed. This process is repeated for every decision time, until the moving window reaches the end of the contract period. The vector of current states $x_{0:K-1}$ and the associated vector of SMPC releases $u_{0:K-1}$ implicitly define a set of time-dependent decision functions $\mu_{0:K-1}$ through the relationship $u_k = \mu_k(x_k)$ for $k = 0 : K - 1$. For convenience, we refer to the SMPC decision as a function in the discussion below, even though SMPC does not explicitly derive such a function.

The distinctive aspects of model predictive control include:

- releases are evaluated only for observed state values, not all possible values;
- the decision function in our formulation depends on the contract energy;
- the decision function is defined implicitly and is available during operations only at the current time (not earlier);

- future revenue is evaluated approximately, over a limited duration time horizon;
- performance is dependent on the accuracy of the predictive inflow and storage models (the stochastic state equations)
- SMPC is approximate, even in the limit as the time horizon becomes infinitely long, because it does not account for the impact of the future measurements,
- computational effort is generally less than SDP, especially for large problems.

2.3.3 Standard operating policy

Both SDP and SMPC make an effort to predict the effect of future uncertain performance by averaging present value revenue over an ensemble of possible inflow disturbances. By contrast, deterministic (non-predictive) operating rules, such as the Standard Operating Policy (SOP) ([78],[82]) do not consider the possible impact of future inflows. These rules typically are heuristic and time-invariant (Figure 2-4). They do not optimize a particular objective and they are specified rather than derived functions of the system state. Non-predictive standard operating policies are easy to implement and convenient for multi-purpose reservoir operations but cannot generally be expected to perform as well in a single-purpose hydropower application as alternatives that utilize information about inflow variability and reservoir dynamics. They are considered here because they are widely used in practice and they provide benchmarks for assessing the potential performance improvement offered by predictive operating rules such as SDP and SMPC. Figure 2-4 shows two SOP variants. The simplest option, indicated by the black curve, releases all available water up to a nominal value equal to the mean inflow $u_{nom} = \bar{I}$ when the storage is $S_{nom} - 0.1\Delta S_{max}$. This nominal release is maintained until a nominal storage level $S_{nom} + 0.1\Delta S_{max}$ is reached. At that point additional water is released up to the maximum turbine capacity u_{max} . Beyond that, excess water must be spilled. The modified red curve hedges the release rule by smoothing abrupt transitions between low, nominal, and high storage conditions.

2.3.4 Perfect information

Reservoir releases and revenues derived by assuming perfect knowledge of future inflows provide useful upper bounds on the performance that can be obtained for a particular actual inflow. In this case releases can be expressed in terms of a decision function but they need not be derived in real time. Instead, they can be computed by maximizing 2.15 with the assumption that the inflow disturbances $\omega_{0:i-1} = \omega^a_{0:i-1}$ are not random but are known perfectly:

$$J_{PI}(u_{0:K-1}, \omega_{0:K-1}, x^a_0, E_c) = \sum_{i=0}^{K-1}(1+r)^{-i}[g[E_i(x^a_0, u_{0:i}, \omega^a_{0:i}), E_c] - \alpha_Z Z_i(x^a_0, u_{0:i}, \omega^a_{0:i})] \\ g_K(x^a_0, u_{0:K-1}, \omega^a_{0:K-1}) \tag{2.23}$$

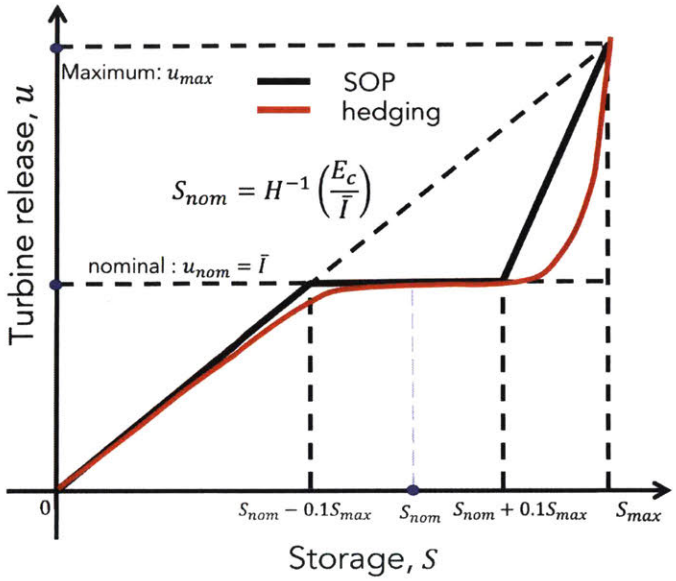

Figure 2-4: Schematic representation of two typical Standard Operating Policies, with the reservoir release expressed as a function of currently available storage. Deviations of the red (hedged) curve from the black (standard) curve indicate an effort to moderate abrupt transitions between low, nominal, and high storage conditions.

This problem can be solved with a standard non-linear programming algorithm since perfect information allows all releases to be computed at once, in batch rather than real-time mode. No reservoir operations method with imperfect information can do better than the perfect information case when presented with the same actual inflow.

2.4 Results and discussion

2.4.1 Setup of the example problem

The problem formulation and solution methods described above are tested here on a typical example using an ensemble of synthetically generated inflows. This Monte Carlo approach enables us to derive revenue probability distributions that quantify the risk associated with different contract selection/ real-time operations strategies. We suppose that the reservoir is designed primarily to generate hydropower and has characteristics similar to facilities such as Hoover Dam, USA; Tehri Dam, India; or and Itaipu Dam, Paraguay ([6],[22]). Figure 2-5 shows the generic reservoir geometry and head vs. storage relation used in our example. The methods of this paper can be applied to any reservoir geometry as long as the storage vs. surface area and the head functions are provided. The standard operating policy used in the example is a non-dimensional version of Figure 2-4 that provides hedging by using a cubic function

(red curve) to smooth transitions between the straight lines (black curve). The black lines are defined by the storage and release break points indicated in the figure.

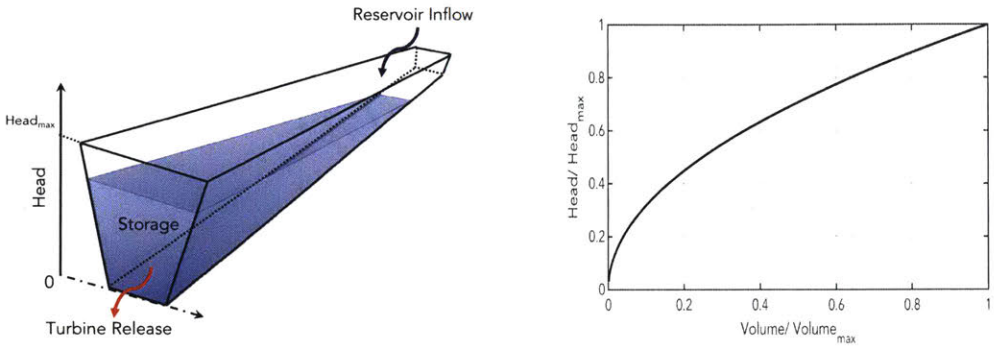

Figure 2-5: Reservoir geometry for the example problem. Left panel shows reservoir configuration and right panel plots the storage vs. head curve for the example

For the example we consider a single state random inflow model that gives sufficient variability to examine firm power shortages and surpluses as well as occasional spills. The normalized log of the inflow is a positive AR1 time series generated from a specified mean inflow, variance, and single lag correlation. The corresponding state equations are special cases of 2.14 and 2.6.

$$\begin{aligned} S_{k+1} &= f_S(S_k, \psi_k, u_k, \omega_k) \\ &= S_k + \Delta t[I_{k+1} - u_k] - Z_k \\ &= S_k + \Delta t[\bar{I}\exp(\rho_\psi \psi_k + \omega_k) - u_k] - Z_k \\ \psi_{k+1} &= f_\psi(\psi_k, \omega_k) = \rho_\psi \psi_k + \omega_k \\ & \psi_0 \sim \mathcal{N}(\bar{\psi}, \sigma_\psi^2) \ ; \ \omega_k \sim \mathcal{N}(\bar{\omega}, \sigma_\omega^2) \end{aligned} \tag{2.24}$$

where $\rho_p si$ is the single lag correlation of ψ_k and the log normal inflow I_k is related to the unitless state ψ_k by:

$$I_k = M(\psi_k) = \bar{I}\exp\psi_k \tag{2.25}$$

The time-invariant mean and variance of ψ_k are computed from the specified mean and variance of I_k :

$$\bar{\psi} = -\frac{\sigma_\psi^2}{2}; \quad \sigma_\psi^2 = \ln(\frac{\sigma_I^2}{\bar{I}^2} + 1) \tag{2.26}$$

$\bar{I}, \sigma_I^2$ are specified inflow mean and variance. The time-invariant mean and variance of ω_k are obtained from:

$$\bar{\omega} = (1-\rho)\bar{\psi}; \quad \sigma_\omega^2 = (1-\rho^2)\sigma_\psi^2 \tag{2.27}$$

The AR1 model has the advantage, for testing purposes, of being less smooth

than higher-order autoregressive models but a more realistic representation of inflows to a moderate size reservoir than noisier time series models. Seasonality could be readily added if appropriate. In practice, the time series model should be estimated from historical inflow data and should be kept sufficiently low-dimensional to make an ensemble analysis of the predictive decision strategies computationally feasible.

A sensitivity analysis of the results can be conveniently formulated in terms of a limited of non-dimensional variables and inputs that are formed from groups of dimensional variables introduced above, using the definitions given in Tables 2.1 and 2.2. These non-dimensional quantities are identified by primed subscripts. Note that there is no spill penalty ($\alpha_Z = 0$) in the nominal case. Also, the maximum possible sustainable energy $E_{max} = \phi(\bar{I}, h_{max}, h_{max})$ appearing in Table 2.1 is achieved when the reservoir head is fixed at its maximum value $h_{max} = H(S_{max})$ and the reservoir inflow and turbine release are both fixed at $\bar{I}$. The actual energy generated over a given time step could exceed this value if the release exceeds the mean inflow. For the example the dimensional problem objective function given in equation 2.13 and the dimensional constraints given in 2.5 through 2.12 are converted to non-dimensional forms by applying the definitions in Tables 2.1 and 2.2 as described in Appendix A. All plots and sensitivity analysis results are expressed in terms of non-dimensional variables.

Non-dimensional variable	Definition	Range or distribution
Storage	$S'_k = \frac{S_k}{S_{max}}$	0.0-1.0
Head	$h'_k = \frac{h_k}{h_{max}}$	0.0-1.0
Inflow	$I'_k = \frac{I_k}{\bar{I}}$	Lognormal
Log Inflow	$\psi'_k = \log I'_k$	0.0-1.0
Release	$u'_k = \frac{u_k}{\bar{I}}$	Non-negative
Spill	$Z'_k = \frac{Z_k}{S_{max}}$	Non-negative
Current Revenue	$g'_k = \frac{g_k}{\alpha_c E_{max}}$	Non-negative
Energy	$E'_k = \frac{E_k}{E_{max}}$; $E'_c = \frac{E_c}{E_{max}}$	0.0-1.0
Objective	$R = \frac{J'_{\mu_{0:K-1}}}{\alpha_c E_{max}}$	Non-negative

Table 2.1: Non dimensional variables

The following subsections examine the results obtained by simulating the reservoir operation with four different coupled contract selection / real-time operations strategies based on Stochastic Dynamic Programming (SDP), Stochastic model Predictive

Non-dimensional input	Definition	Value in example
Reservoir residence time	$\tau_{res} = \frac{S_{max}}{I\Delta t}$	Nominal $\tau_{res}^{low} = 12$; $\tau_{res}^{high} = 48$
Maximum reservoir release	$u'_{max} = \frac{u_{max}}{I}$	1.5
Contract Period	$K' = \frac{K}{\Delta t}$	100
MPC window length	$w' = \frac{w}{\Delta t}$	12
Spill penalty coefficient	$\alpha'_z = \frac{\alpha_z S_{max}}{\alpha_c E_{max}}$	Nominal: $\alpha_z^{low} = 0$; $\alpha_z^{high} = 20$
Revenue Coefficients	$\alpha'_1 = \frac{\alpha_1}{\alpha_c}$; $\alpha'_2 = \frac{\alpha_2}{\alpha_c}$	$\alpha'_1 = 2$; $\alpha'_2 = 0.15$
Log inflow parameters	ρ_ψ, σ_ψ^2	$\rho_\psi = 0.8$, $\sigma_\psi^2 = 0.18$
Number of replicates	N	50
Number of meta-replicates	N^a	200
Discount factor	r	4%

Table 2.2: Non dimensional inputs

Control (SMPC), a Standard Operating Policy (SOP) and a Perfect Information Scenario (PIS). They also consider the effect of varying influential dimensionless inputs such as the non-dimensional residence time, spill coefficient, and log inflow statistics.

2.4.2 Hydropower revenue comparison

The overall performance of the four decision strategies described in Section 2.3 can be assessed in terms of a number of performance measures, such as the net present value of the hydropower revenue generated over the contract period, revenue volatility over time, spill magnitude and frequency, etc. In our ensemble analysis many of these performance measures are random variables by virtue of their dependence on random inflows. To illustrate the capabilities of an ensemble approach we compare probability distributions for the net present value of the four decision strategies introduced earlier. Similar comparisons can be made of other performance measures. It is convenient to compare revenue performance in terms of the dimensionless revenue ratio R defined in Table 2.1. We first consider performance for the nominal input values given in Table 2.2 and then for a few alternatives that use different values for some of these inputs.

The perfect information strategy is unique among those considered here since it relies on advance knowledge of the entire sequence of reservoir inflows. With perfect inflow information, it is possible to derive a different optimum E_c for each meta-replicate in the Monte Carlo simulation. By contrast, each of the other strategies

work with a single E_c value that maximizes expected revenue over the entire inflow ensemble for that particular strategy.

Figure 2-6 compares the kernel density estimates of probability distribution of the revenue ratio for all four decision strategies for nominal inputs. The variation in revenue observed for the perfect information (PIS) case depends only on the intrinsic variability of the actual inflow, not on the algorithm's ability to predict this inflow (since it has access to perfect inflow information). If the inflow for a particular actual inflow meta-replicate is low for a prolonged period, revenue will be low, even though the inflow is known perfectly. The other three decision rules are affected both by the intrinsic variability of the actual inflow and by uncertainty in the inflow predictions used to make release decisions. That is why there distributions are shifted to the left, toward lower revenue. The SDP and SMPC strategies tend to be more sharply peaked near their modes but have relatively long tails at lower revenue values, reflecting the consequences of occasional poor predictions. The most visible property of the PIS is its greater probability of yielding high revenue ($R > 0.75$).

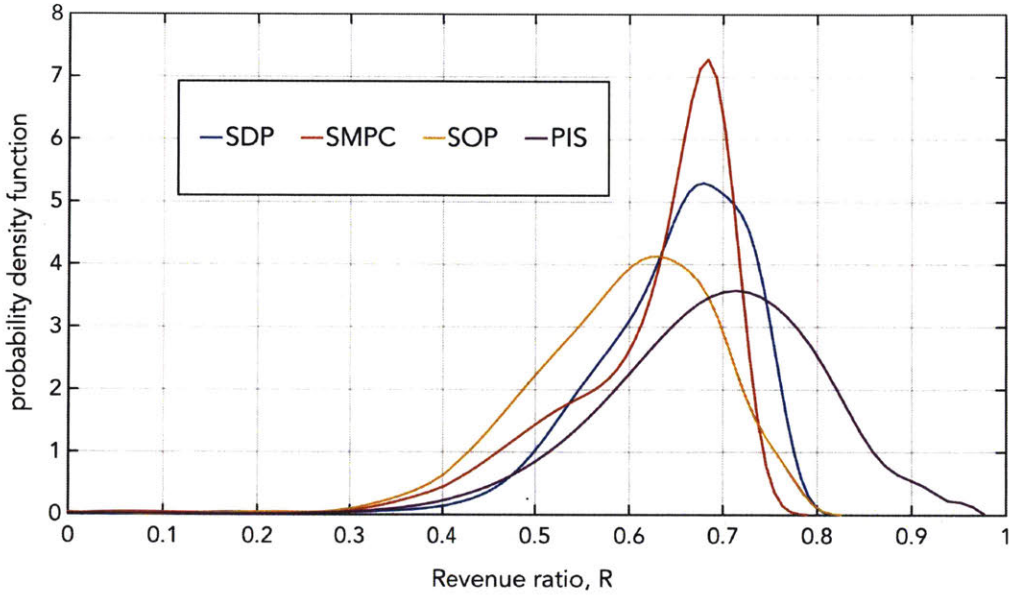

Figure 2-6: Probability density function of the revenue ratio for SDP, SMPC, SOP and PIS operational techniques

Stochastic dynamic programming (SDP) is second among the alternatives in terms of mean revenue since it makes best use of the ensemble inflow predictions when optimizing the current release. The backward recursion stores release strategies that maximize the expected revenue for the remaining contract time from any value of the state. These strategies can be recovered as the actual state values become known. By contrast, stochastic model predictive control (SMPC) derives a current release that maximizes expected revenue only from the current state. The replicates used in this calculation may not reflect the actual evolution of the system at later times. Also, the SMPC maximization is limited to a window that can be significantly shorter than

the remaining contract time. For these reasons, SMPC is somewhat less likely to give high revenues and more likely to give low revenues than SDP [42]. The non-predictive standard operating policy performs the worst among the four alternatives, generating the smallest mean revenue with the highest probability of low revenues. This reflects the method's inability to adjust releases when near-future inflows and storages are likely to be lower or higher than average, given current inflow and storage. By contrast, predictive methods such as SDP and SMPC adjust releases in anticipation of possible future conditions. Table 2.3 lists the average revenue ratio computed over all the inflow meta-replicates as well as the probability (in %) of achieving a low revenue ratio below 0.5 or high ratio above 0.75. These percentages complement information on the mean revenue by considering the probability of low or high revenue values when comparing decision strategies.

2.4.3 Sample time series

The Monte Carlo simulation conducted in our example provides individual replicates of relevant dynamic variables such as the inflow, storage, release, and energy output as well as the revenue probability distributions discussed above. Figure 2-7 compares these variables for four different decision strategies, all using the nominal inputs from Table 2.2. Each of these four cases maximizes one of the decision strategy objectives specified in Section 2.3 by selecting the best possible combination of contract firm energy and release history for a given actual inflow meta-replicate. The normalized values of E_c for this example (expressed as a fraction of E_{max}) are 0.61 for perfect information (PIS), 0.57 for stochastic dynamic programming (SDP), 0.51 for stochastic model predictive control (SMPC), and 0.48 for the standard operating policy (SOP). Comparing to Figure 2-6, the predictive strategies that generate higher firm power are also more likely to produce higher revenue.

The top panel of Figure 2-7 shows the non-dimensional reservoir inflow series together with four turbine release series computed in real time from the current storage and inflow values, one for each of the four operating rules. The middle panel shows the non-dimensional reservoir storage generated from these releases, with the maximum normalized storage given by 1.0. The challenge for the operating rule is to keep water levels high in order to maximize energy output while avoiding spills that may have adverse downstream consequences and that also reduce the quantity of water available for generating power.

In the nominal case shown in Figure 2-7 the SDP decision strategy generally maintains higher storage than the other techniques, often approaching the reservoir capacity. This reflects SDP's somewhat better predictive capabilities and also the fact that spills are not explicitly penalized in the nominal case. SMPC behaves similarly but gives somewhat more erratic releases and energy production. Higher variability in energy together with a somewhat lower firm power value yield somewhat lower revenue for SMPC. The PIS is able to maintain much more stable release and energy production levels than any of the other methods. This reflects its ability to adjust releases in anticipation of future high or low inflow events, which are known perfectly. The advantage of perfect information also allows PIS to maintain a storage

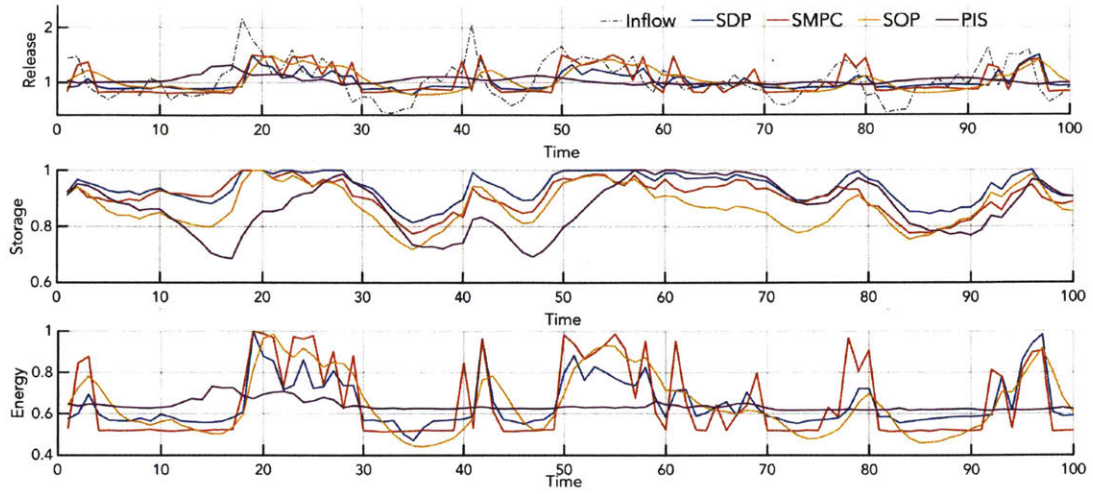

Figure 2-7: Example reservoir operations with the four techniques (SDP, SMPC, SOP and PIS) plotted for a particular inflow meta-replicate. Top panel: Reservoir inflow time series and turbine release; Middle-panel: Reservoir storage; Bottom panel: Energy generated. All quantities are non-dimensional.

level that is generally lower than the other alternatives, even though the PIS average energy production and revenue are higher. The PIS result suggests the level of performance that SDP and SMPC could approach if they had access to very accurate inflow estimates.

2.4.4 Sensitivity analysis

All of the non-dimensional parameters listed in Table 2.2 affect the performance of the four different operational strategies considered here. It is useful to examine in detail two key dimensionless inputs, the normalized spill penalty coefficient $\alpha_z^{'}$ and the residence time τ_{res}, and to briefly consider some of the others.

Sensitivity to spill penalty

A higher spill penalty tends to make the operational strategy more conservative, lowering the water level below the maximum to reduce the magnitude and frequency of spills. Table 2.3 includes a comparison of expected revenue and the probability of low and high revenues for a moderately high non-dimensional spill penalty value vs. the nominal case that does not penalize spills. Figure 2-8 shows revenue ratio probability distributions for the same two spill penalty options. Increasing the spill penalty consistently shifts the revenue probability density towards lower values (Figure 2-8). As the penalty coefficient increases spill occurrences decrease from 15% to 2.7% for dynamic programming, from 7% to 2.1% for model predictive control, and from 6% to 2.8% for the standard operating policy. Dynamic programming has the highest spill occurrence for the unpenalized case because its more complete description of uncertain

future conditions benefits more from pushing the reservoir system to capacity in order to achieve maximum performance. Its more complete treatment of uncertainty also enables dynamic programming to significantly reduce spill occurrence when spills are penalized. By contrast, SOP gives a lower unpenalized spill occurrence but does not achieve as great a reduction when spills are penalized. Model predictive control falls somewhere in between.

The perfect information option shows a similar sensitivity to the spill penalty but gives a lower unpenalized spill occurrence than any of the alternatives. Perfect information makes the most difference during high inflow events that can cause spills since it enables the operating rule to draw down the reservoir before high flows occur. By reducing the amount of water lost to spills the perfect information option is able to generate more hydropower and greater revenue. It is possible to decrease spill occurrence somewhat further than indicated in Figure 2-8, by further increasing the spill penalty. But this effect is ultimately limited by the inflow statistics. Overall, perfect information and dynamic programming sacrifice revenue less than the other alternatives when spills are penalized.

Technique	Low spill penalty $\alpha_z^{low} = 0$ *(nominal)*			High spill penalty $\alpha_z^{high} = 20$		
	E(R)	$P(R < 0.5)$	$P(R > 0.75)$	E(R)	$P(R < 0.5)$	$P(R > 0.75)$
Perfect Information	0.69	3%	25%	0.62	18%	27%
Dynamic Programming	0.64	5%	8%	0.62	10%	4%
Model Predictive control	0.62	11%	6%	0.58	26%	2%
Standard operating	0.59	20%	5%	0.51	40%	0%

Table 2.3: Comparison of the average revenue ratio R and probability of low R (< 0.5) and a high R (>0.75) between the four operational strategies

Sensitivity to residence time

The residence time (τ_{res}) provides a concise description of the combined effect of the reservoir capacity and the mean inflow. Increasing the residence time (low inflows/ large reservoirs) reduces sensitivity to inflow variability, generating higher revenue for extended periods. Figure 2-9 shows this behavior by plotting the revenue ratio probability distributions for two different residence time options: low τ_{res}^{low} (nominal) vs. high τ_{res}^{high}. Increasing the residence time reduces the effects of inflow variability and shifts the revenue distributions towards higher values consistently across all four techniques. The revenue distribution also narrows, reducing the risk of lower revenues. Increasing the residence time increases the average revenue by 18% for dynamic programming, 20% for model predictive control, and 18% for the standard operating policy.

With a high residence time, high inflow events do not necessarily cause uncontrolled spills. They can be captured as storage, making it possible to temporarily

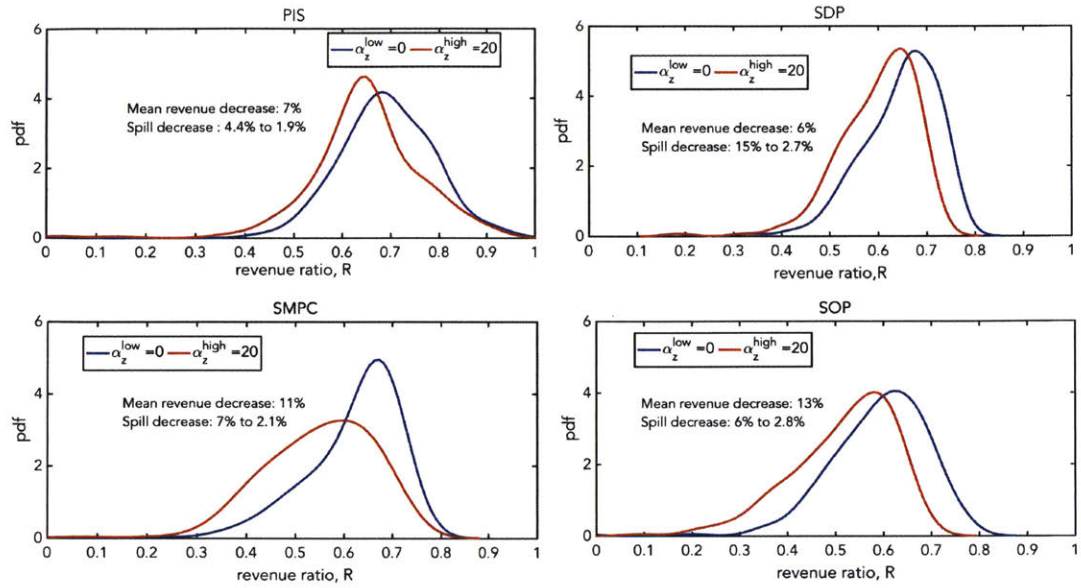

Figure 2-8: Effect of spill penalty on the revenue density function. Mean revenue and spill frequency both decrease as the spill penalty is increased from nominal $\alpha_z^{low} = 0$ to $\alpha_z^{high} = 20$

allow releases greater than the mean inflow ($\bar{I}$). This can yield revenue ratios (R) greater than 1 (see, for example, the perfect information case in Figure 2-9). For the nominal spill penalty coefficient spill occurrences decrease significantly with increasing residence time: 4.4% to 0.15% for perfect information, 15% to 3.8% for dynamic programming, 7% to 1.9% for model predictive control and 6% to 0.3% for standard operating policy. Although a large residence time reservoir is clearly desirable the potential for increased capacity is practically limited by site constraints and higher costs. When designing a new reservoir such considerations need to be included in the optimization process.

Technique	Low residence time $\tau_{res}^{low} = 12$ *(nominal)*			High residence time $\tau_{res}^{high} = 20$		
	E(R)	$P(R < 0.5)$	$P(R > 0.75)$	E(R)	$P(R < 0.5)$	$P(R > 0.75)$
Perfect Information	0.69	3%	25%	0.82	0%	75%
Dynamic Programming	0.64	5%	8%	0.78	1%	79%
Model Predictive control	0.62	11%	6%	0.77	3%	78%
Standard operating	0.59	20%	5%	0.70	6%	38%

Table 2.4: Comparison of the average revenue ratio R and probability of low R ($<$ 0.5) and a high R ($>$0.75) between the four operational strategies

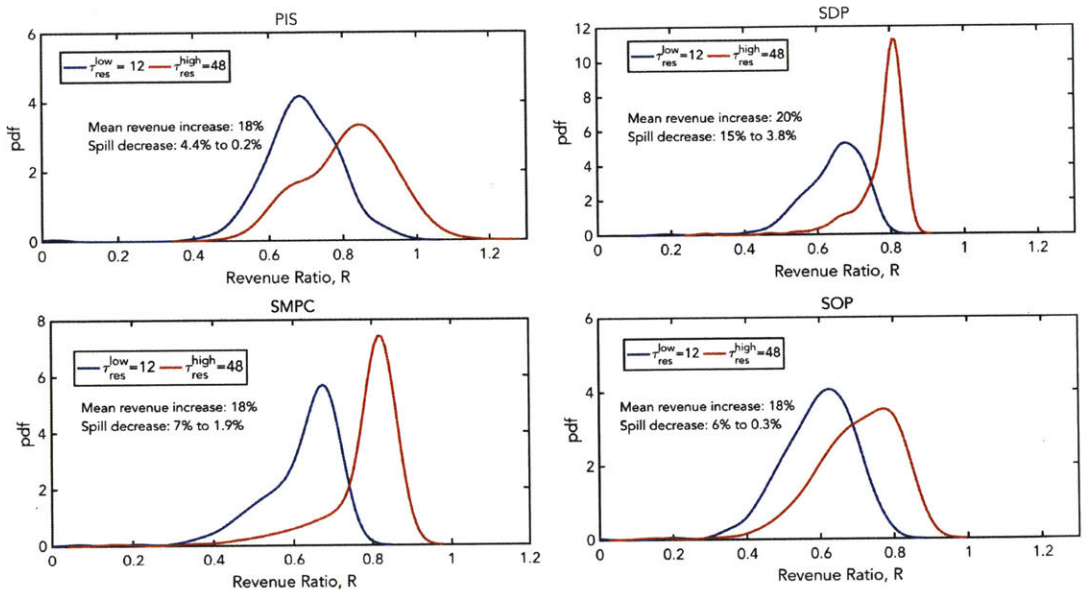

Figure 2-9: Effect of residence time on the revenue density function. Increase in residence time τ_{res} (nominal $\tau_{res}^{low} = 12$, $\tau_{res}^{high} = 48$) shifts the revenue distributions to higher revenue in every operational strategy

Sensitivity to other factors

The preceding sections show that the comparative performance between the three real-time operational strategies (SDP, SMPC and SOP) is sensitive to spill penalty and residence time. Performance also depends on other parameters such as reservoir geometry, inflow statistics, discount rate, and revenue function coefficients. For example, predictive operating strategies such as SDP and SMPC provide a greater performance benefit if the reservoir inflow series has a high serial correlation ρ_ψ and a low or moderate variance σ_ψ^2. In such cases it is easier to predict near-term inflows. On the other hand, if the correlation is close to 1 and the variance is high the possibility of extended periods of anomalous inflows leads to reduced benefit even for predictive algorithms. The revenue function parameters can also influence performance through their impact on both the contract value and real-time operations. For example, increasing α_1 increases the penalty of generating a shortfall, which leads to a more conservative contract that keeps reservoir storage near capacity and increases spill occurrence. The combined effect of many sensitivities determines the relative effectiveness of predictive vs. deterministic operating rules in any given situation.

2.5 Conclusions

This study prescribes a novel stochastic optimization approach that simultaneously handles contract energy negotiation and real-time control of a hydropower reservoir. Contract selection and revenue performance are closely coupled to the choice

of operational strategy used. Predictive techniques such as stochastic dynamic programming (SDP) and stochastic model predictive control (SMPC) give significantly better revenue (mean improvement $> 10\%$) than a non-predictive standard operating policy (SOP) for the nominal conditions considered here. For other conditions the improvement may be either greater or less. Predictive techniques tend to work best in situations where reservoir inflow statistics favor the use of inflow and storage forecasts for optimizing revenue. Between the two predictive techniques, SDP generates higher revenue than SMPC but can be more computationally demanding, especially for multi-reservoir systems.

Sensitivity analysis indicates that a high spill penalty has a negative impact on revenue since it leads to strategies that operate the reservoir at a lower storage level. Reservoirs with a higher residence time generate higher revenues and result in less spill since the sensitivity to inflow variability decreases.

In addition to focusing on firm power deliveries and revenue, the conceptual framework presented here provides a probabilistic perspective that quantifies both revenue and spill risk. This framework can be readily adapted to accommodate different reservoir shapes, inflow statistics and revenue functions. It can also be extended to multi-reservoir systems. A stochastic approach that focuses on the probability distribution of revenue can provide useful insights and tangible benefits for both hydropower reservoir operations and contract negotiations.

Chapter 3

Managing groundwater as a common pool resource: A multi-compartmental approach

3.1 Introduction and context

3.1.1 Background

Groundwater is the largest reserve of liquid fresh water on earth. It plays an especially important role in providing a reliable water supply for irrigated agriculture, contributing about 40% of all irrigation water ([2],[31],[69]). Given this importance, it is natural that there is widespread concern about declines in groundwater storage in the western United States, northern India, northeastern China, and other important farming regions. Although groundwater storage naturally fluctuates with variations in local climate most of the decline in regions with significant depletion is due to human activities, particularly increases in agricultural withdrawals ([38],[77],[67]).

Since a groundwater reservoir acts as a buffer that smooths out short-term meteorological fluctuations, some intentional variation in storage can be considered to be good management practice. However, when aquifer outflows remain above inflows for an extended period of time groundwater levels can decline sufficiently to have adverse impacts. These include increases in cost and decreases in accessibility as well as undesirable impacts on groundwater water quality if drawdowns increase the intrusion of water from formations with higher salt content. Groundwater depletion can also decrease recharge to surface water ecosystems that are closely coupled to groundwater. The effects of depletion can be irreversible in cases where drawdowns are sufficient to change aquifer storage capacity through subsidence ([25],[21])

The risks of long-term groundwater depletion have focused considerable attention on better understanding of the factors that may encourage or discourage undesirable depletion. Many researchers have posed groundwater depletion as a common pool problem and, more specifically, as a non-cooperative non-zero sum dynamic game ([54],[59],[63],[51],[49]). The resource is characterized as non-exclusive and openly accessible to a number of players (or agents). Each player competes to acquire, over

time, as much of the resource as needed to maximize his own discounted net benefit (revenue minus cost), subject to physical constraints.

There is an incentive for each player to remove groundwater sooner, when water levels are higher and pumping is less expensive, than later, after water levels have dropped in response to the collective actions of all the players. The situation is further complicated by the possibility that a player's access to groundwater may be lost when the local water level falls below the well depth. The incentive to extract groundwater before water levels fall further can lower net benefit for all unless mechanisms for cooperation are introduced. This result is supported by game theoretic analysis, which provides a way to evaluate the strategic externality that results from uncooperative behavior ([54],[59],[63],[64]). However, the magnitude of this externaltity and the significance of the economic inefficiency induced by uncooperative groundwater exploitation have been debated [39].

A number of studies have analyzed common pool water management as a single player optimization problem. In this case the problem is posed from either i) the perspective of a *'social planner'* who manages the entire resource or directs everyone's individual water use to maximize an appropriate measure of aggregate long-term benefit or ii) the perspective of a *'myopic'* player pumping without regard for future consequences or the actions of other players. ([81],[68],[29]). In the context of groundwater depletion, the first option is typically solved as an optimal control problem. The second option is solved by equating the current marginal benefit of pumping to the current marginal cost. Gisser [3] compare these extremes using a simplified *'bucket model'* of a groundwater aquifer. When all pumped water is evapotranspired this model can be written in discrete time as:

$$\frac{\Delta S_k}{\Delta t} = A\theta \frac{\Delta h_k}{\Delta t} = R - \sum_i u_k^i \tag{3.1}$$

Where ΔS is the change in the stock over stage (or decision interval) k. The time index identifies $k = 1, , K$ stages that are each of duration Δt , with stage k extending from t_k to t_{k+1} . The specified input A is the horizontal area of the study region, θ is the specific yield for unconfined aquifers and the storage coefficient for confined aquifers, and h_k is the spatially averaged hydraulic head. The decision or control variable u_k^i is the player i's pumping rate over stage k, determined at time t_k . Equation 3.1 is initialized with a specified head value h_1 . The only external source of water in the bucket model is a constant recharge R that does not depend on pumping. Also, there is no lower limit on storage so the only factor restricting aquifer depletion is pumping cost.

The Gisser and Sanchez [3] analysis gives nearly the same exponentially decaying pumping rates and groundwater levels for the social optimal and myopic options. This result has been widely interpreted to contradict the presumption that a player pursuing only short-term gain will deplete a groundwater resource much more than a player with a longer-term perspective. Koundouri [39] provides a comprehensive review of research prompted by the Gisser and Sanchez result as well as of broader issues related to the control of groundwater depletion. She indicates that a number of

case studies support the Gisser and Sanchez conclusions. Other relevant assessments of the single-player management problem include Loaiciga [45], Brozovic [14], and Madani [46]. None of these directly address non-cooperative behavior of multiple players.

Studies of the multi-player non-cooperative groundwater depletion problem often focus primarily on qualitative properties of Nash equilibrium solutions ([54],[59]). Some game theoretic analyses have been able to obtain quantitative solutions by making various simplifications, including steady-state assumptions ([63],[64],[51]). Nakao et. al [51] quantify the strategic effects of non-cooperation using a bucket model for the Hueco Bolson aquifer shared by Texas and Mexico. Mueller et al [49] solved a two time-step dynamic groundwater game using time-averaged hydrologic response functions derived from a fully dynamic and spatially distributed groundwater model. This paper builds on related studies by Saak and Peterson [65] and Saleh et. al [66] that account for the spatial configuration of the player's wells.

Multi-player groundwater game studies have clarified important distinctions between open loop (path) and closed loop (decision rule) Nash equilibria and have identified the roles of pumping cost and storage limits in restraining depletion. However, they either do not quantify the magnitude of the inefficiency caused by non-cooperative behavior or they look only at specialized situations that permit the game problem to be readily solved. In order to use game theory to better understand the causes and implications of non-cooperative depletion we need a problem description that captures essential economic and hydrologic features while also remaining computationally tractable.

3.1.2 Factors influencing the realism of groundwater common pool formulations

Benefit and Cost functions

Since groundwater pumped for irrigated agriculture is typically used on site to grow crops its value is measured in terms of the derived demand for agricultural products. A particular player's demand for groundwater can change in response to changes in water availability from other sources (including rainfall) or changes in the market prices and/or the yields of alternative crops. Groundwater extraction is affected by energy costs as well as by the behavior of other players. Uncertain fluctuations in meteorological variables, crop prices, and the cost of energy imply a need to formulate the problem in terms of real-time closed loop (or feedback) rather than open-loop player decision rules since benefits and costs can be expected to change over time.

Boundary conditions and steady-state solutions

The bucket model of equation 3.1 is the simplest possible approximation of a realistic groundwater model. Its single state is the bucket water level, which can be interpreted as a spatially averaged value of hydraulic head. The only source of water is a specified constant recharge and water can leave only through pumping ([54],[37],[65],[74],[57]).

These assumptions imply that the bucket model is at steady state only if the total water pumped is exactly equal to the constant recharge. Also, pumping costs are the same for all players and depend on the spatially averaged head rather than the drawdown at individual player wells. Bredehoeft [13] points out the limitations of sustainability assessments based on the bucket model. In particular, he notes that specified head aquifer boundaries often play an important role in practice and need to be included. The time-dependent fluxes crossing such boundaries generally depend on the system state as well as the specified boundary heads.

Stock and boundary flux limitations

In most game theoretic analyses of the groundwater common pool problem pumping cost is the only factor limiting extraction. This is because the aquifer depth, and therefore the stock of groundwater, is assumed to be infinite, so that the aquifer can be depleted indefinitely at an arbitrary rate except for restraints imposed by pumping cost. In reality, continuous pumping at sufficiently high rates can eventually deplete the groundwater stock or reduce boundary fluxes from external sources such as rivers, putting physical limits on pumping even when cost restraints do not apply. Stock and flux limitations can be important when the pumping cost is low (e.g. subsidized) and/or the aquifer depth is relatively small. However, these effects are rarely included in non-cooperative game formulations because of the mathematical complications they introduce. Stock and boundary flux limitations need to be considered to provide for the full range of possibilities in aquifer management.

Player heterogeneity and number of players

The realism of a game description of the common pool groundwater problem is also affected by assumptions made about player heterogeneity [65]. Although it may be appropriate to assume that all agents are identical in some situations it is important to understand how player differences can influence individual behavior and the overall impact of exploitation when they do occur. Of particular interest are differences that arise because different players grow different crops with different prices, yields, and water requirements. These differences affect each agent's pumping decisions, with some agents being rewarded more for increased pumping than others, even though they share a common resource. An additional related factor is the effect of increasing the number of players. It is particularly helpful to know how revenue efficiency might change as the number of players increases above two.

In this chapter we consider a multi-compartment stream-aquifer model that addresses most of the issues identified above. Our approach expands the scope of classical Nash equilibrium solutions to the groundwater common pool problem while retaining the capability to derive closed form decision rules that apply for the important special case of a quadratic objective with linear equality constraints ([7],[24],[1]). The existence of a convenient closed form Nash equilibrium solution enables us consider a number of interesting phenomena, including aquifer-river interactions, asymmetric agent benefit functions, and the effects of dynamic changes in recharge and other

inputs. Taken together, these investigations provide useful insight about player behavior over a range of decision strategies, hydrologic and economic conditions.

3.2 Formulation of the groundwater common pool problem

3.2.1 Objectives and constraints

The groundwater common pool problem can be posed as an N-person finite horizon non-cooperative discrete time deterministic game, following notation used by Basar and Oldser [8]. The objective functional to be maximized by Player i is:

$$\begin{aligned} L^i(u^1,...,u^N) &= \sum_{k=1}^{K-1} g_k^i(u_k^1,...,u_k^N,x_k,x_{k+1}) + g_K(x_K); \quad i=1,...,N \\ &= \sum_{k=1}^{K-1}(1+r)^{-k}\left[B_k^i(u_k^i) - C_k^i(u_k^i,x_k)\right] + g_K(x_K) \end{aligned} \tag{3.2}$$

where u_k^i is the control variable (water pumped, L^3T^{-1}) defined for Player i over stage k and is a vector of Player i pumping rates for all K stages. The function g_k^i is the net benefit obtained by Player i over time step (or stage) k, is the system state vector at discrete time t_k , and $g_K(x_K)$ is a specified terminal net benefit intended to account for residual value of the resource at the end of the finite horizon. The discounted net benefit in each stage($) is proportional to the difference between the current revenue $B_k^i(u_k^i)$ from using pumped groundwater to grow crops and the current cost of pumping $C_k^i(u_k^i,x_k)$, where r is the discount rate. We assume that land is not a limiting resource for any of the players and that the revenue function for each player is strictly concave and monotonically increasing. Note that the current revenue for player i depends only on the water that player pumps, not on the state. Also, the current cost depends, via the state, on the past and present pumping rates of all players.

The state vector is determined by a state equation that describes the temporal evolution of the groundwater system. The discrete time form of this equation is:

$$x_{k+1} = f_k(x_k,u_k^1,...,u_k^N); \quad k=1,...,K-1; \quad x_1 \; given \tag{3.3}$$

In the rest of this paper we assume a constant discrete time step of one year to coincide with a typical crop cycle. All flux variables are annual totals and all states are defined at the beginning of the water year. The dynamic game problem defined in equations 3.2 and 3.3 does not have a unique solution unless further restrictions or information are included. One way to narrow the scope of possible solutions is to seek Nash equilibria [8], as discussed in the next section. Another is to consider only a single player.

In principle, the state equation for the groundwater common pool problem can

be derived by discretizing the spatially distributed groundwater flow equation and associated initial and boundary conditions over time and space. This approach can account for the complexities of real world aquifers, including local drawdowns at pumping wells, but it generates a large state vector that is comparable in size to the number of discrete grid cells and greatly increases the computational difficulty of the dynamic game problem. An alternative popular in classical optimization studies is to use a time-dependent version of a response matrix, which relates the spatially averaged heads in pumping cells directly to the player pumping rates. The response matrix can be derived from the discretized groundwater equation, outside the optimization algorithm ([30],[49]). This approach reduces the size of the system state but the resulting state constraint does not have the recursive form given in equation 3.3. This complicates analysis of uncooperative games that rely on dynamic programming methods to derive subgame perfect Nash equilibrium solutions.

Many of the key physical phenomena included in a detailed groundwater model can be captured in a computationally efficient compartmental model that can be viewed as an extension of the bucket model. We demonstrate the concept here with a two-compartment model that distinguishes outer (recharge) and inner (pumping) portions of an unconfined aquifer and also adds a specified head river boundary that can act as either a source or a sink of water. This model, which is illustrated in Figure 3-1, has the following form:

$$A_O\theta_O\frac{\Delta h_{Ok}}{\Delta t} = A_O\theta_O f_{Ok}(x_k) = R_k - \lambda_{OI}(h_{Ok} - h_{Ik}) \tag{3.4}$$

$$A_O\theta_O\frac{\Delta h_{Ok}}{\Delta t} = A_O\theta_O f_{Ok}(x_k) = R_k - \lambda_{OI}(h_{Ok} - h_{Ik}) \tag{3.5}$$

$$x_k = [h_{Ok},\ h_{Ik}]^T \tag{3.6}$$

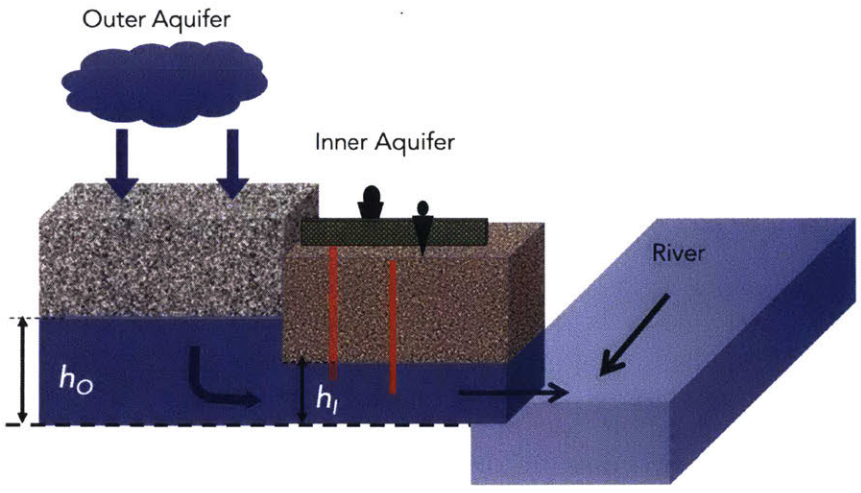

Figure 3-1: Schematic representation of the lumped two-state model

Under natural unpumped conditions groundwater flows from the outer recharge area through the inner aquifer to the river. When pumping is added some of the recharge is diverted from the river for irrigation. If the pumping rate is high enough to draw the inner aquifer head below the river head the flow to the river can reverse, with the river acting as a source of water rather than a sink. The two-compartment model reduces to the bucket model if λ_{OI} becomes very large so $h_{Ik} \simeq h_{Ok}$ and if $\lambda_{IR} = 0$. Then the inner and outer compartment equations can be summed to give a single bucket state equation.

Our analysis of the common pool problem for the multi-compartment model adopts a quadratic current benefit function that reflects a linearly decreasing marginal demand for water ([28],[36]):

$$\begin{aligned} B_k^i(u_k^i) &= b_1^i u_k^i + b_2^i (u_k^i)^2; \quad b_1^i > 0, \; b_2^i < 0 \\ 0 < u_k^i &< u_{max}^i = -\frac{b_1^i}{1 b_2^i} \end{aligned} \tag{3.7}$$

Note that the quadratic benefit function adopted here is monotonically increasing for pumping rates that satisfy the last inequality constraint. This constraint insures that the marginal demand remains non-negative. At the rate u_{max}^i additional water no longer provides any benefit (e.g. because land is limiting).

The current pumping cost for Player i depends on the energy expended to pump groundwater at a rate from the inner aquifer water level to the ground surface elevation h_g :

$$C_k^i(u_k^i, h_{Ik}) = c\eta(h_g - h_{Ik})u_k^i; \quad c_1^i > 0 \tag{3.8}$$

where c represents the price per unit energy consumed and η is the pump efficiency (both of these parameters are assumed to be the same for all players). For simplicity, amortized capital costs of well construction are neglected but they could be included in equation 3.7 if desired. This cost function uses the spatially averaged inner aquifer head h_{Ik} to determine the pumping cost for each of the players. In reality, player pumping costs depend on the local heads at the pumping wells. The local head can be expected to vary throughout the aquifer and will generally be different at different wells. The model's use of the spatially averaged inner aquifer head for cost estimation is an approximation that is more accurate if there are many pumping wells (generally more than one for each player) that are distributed more or less uniformly over the aquifer area and if the aquifer transmissivity is high so the head gradients are small. In that case, the pumped local head will be relatively uniform over the inner aquifer (but lower than the unpumped local head) and the approximation that pumping cost is proportional to h_{Ik} is adequate for investigation of the issues of interest in this paper.

The objectives and state equations given above may be put in an affine-quadratic form if the lateral flow coefficients are constants and the pumping rate constraint in equation 3.7 is inactive. This is particularly convenient for computing Nash equilibrium solutions for discrete time games. The net benefit used here is an extension of

the purely quadratic objective found in most texts on optimal control and dynamic game theory [e.g. Basar and Olsder [8]], since it adds a cross term between the state and control variable as well as terms that are linear in the state and control:

$$L^i(u^1,...,u^N) = {x_K}^T Q^i_{xx,K} x_K + \sum_{k=1}^{K-1} {u^i_k}^T Q^i_{uu,k} u^i_k + {x_k}^T Q^i_{ux,k} u^i_k + S^i_{uk} u^i_k \quad (3.9)$$

$$x_{k+1} = F_k x_k + G^1 u^1_k + ... + G^N u^N_k + E_k \quad (3.10)$$

The particular matrix coefficients can be obtained from the state equation and the objective (3.8,3.7,3.4-3.6) Problems of this form have unique closed form solutions for the decision strategies discussed below. Inequality constraints, such as those needed to impose certain stock limitations, generally require numerical solutions.

3.2.2 Decision Strategies

The objective function and state equations in two-compartment problem description are intentionally selected to give a convenient linear-quadratic solution structure. The theory needed to derive a variety of decision strategies for optimal control and dynamic game problems with this structure is well-established ([7]). In order to apply this theory we need to specify the nature of the decisions to be made by the players. A decision strategy is written as a function $\gamma^i(\cdot)$ that relates the control variable u^i_k for Player i to that player's information about the system at time t_k:

$$u^i_k = \gamma^i_k(\eta^i_k) \quad (3.11)$$

where η^i_K denotes a particular information set that generally includes observations of the system state. Here we will consider four alternative decision strategies.

1. Social optimum strategy
2. Open Loop Nash Equilibrium strategy
3. Subgame perfect or Feedback Nash Equilibrium strategy
4. Myopic strategy

Due to the particular quadratic structure of the objective function and affine structure of the state dynamics all of these decision strategies yield affine decision rules that depend on either the initial or current state:

$$u^i_k = \gamma^i_k(x_1) = \Upsilon^i_k x_1 + \Gamma^i_k \ (\textit{Open Loop}) \quad (3.12)$$

$$u^i_k = \gamma^i_k(x_1) = \Upsilon^i_k x_k + \Gamma^i_k \ (\textit{Social optimum, Feedback, Myopic}) \quad (3.13)$$

Although all four decision rules have an affine form their coefficient matrices have different values and change over time in different ways. This leads to differences in both dynamic and steady-state behavior. The following paragraphs summarize distinctive aspects of the four alternatives.

Social optimum (single player)

This case considers management of the entire aquifer by a single player who selects pumping rates at all wells to maximize the total discounted net benefit obtained from the resource over a finite planning horizon, given perfect knowledge of the current state. This problem can be viewed as equivalent to a cooperative game where multiple players agree to implement the single player solution and then distribute the benefit in a mutually agreeable manner. The social optimum solution is of particular interest because it establishes a benchmark that may be used to assess inefficiencies that arise when multiple players compete in a non-cooperative game. The information set for the single social optimum player consists of the state vector at the current time. The sequence of pumping rates that maximizes 4.10 subject to 4.11 may be found with dynamic programming, following an approach often used to solve a single player optimal control problem. This gives an affine decision rule with coefficient matrices derived from the player objectives and state equations [8]. Details are provided in appendix B.2 towards the end of the study.

Open loop (multiple players with committed pumping decisions)

In an open loop uncooperative game the players maximize their individual discounted net benefits but each commits at the initial time to a particular strategy for the entire planning horizon, given perfect knowledge of the initial state. The open loop strategy is of interest primarily because it is more readily computed than the more realistic closed loop strategy ([20]). Here we seek open loop Nash equilibrium strategies that arise when no player can benefit by changing his/her strategy while other players keep their strategies unchanged. The open loop linear quadratic problem based on 4.10 and 4.11 can be solved as an equivalent static optimization problem for all player pumping rates at all stages or as a coupled set of dynamic optimal control problems, one for each player [8]. The coefficient matrices and conditions for a unique open loop Nash equilibrium are described in the appendix section appendix B.3.

Feedback (multiple players with real-time pumping decisions)

The feedback decision strategy considered here insures that each player maximizes anticipated future benefits at each stage of the game for any value of the current state, using a subgame perfect Nash equilibrium solution. This allows players to respond in real time to unexpected changes in meteorological or cost inputs or to unexpected (perhaps suboptimal) pumping decisions made by other players. A feedback strategy can be less efficient than an open loop strategy since it must exclude options that incur short-term losses to obtain longer term gains ([71],[19],[18]). However, a feedback strategy is more realistic since it recognizes that players will respond dynamically to changing conditions rather than strictly follow a strategy fixed at the initial time. A closed loop subgame perfect Nash equilibrium strategy can be obtained from a coupled set of dynamic programming problems, one for each player, using a backward recursion [8]. For the linear affine problem of equations 4.10 and 4.11 the dynamic programming algorithm takes the form of a set of coupled matrix difference equations.

The coefficient matrices and conditions for a unique feedback Nash equilibrium are described in appendix B.4.

Myopic (multiple players acting without regard for the future)

A player pursuing a myopic strategy maximizes current net benefits only, without regard for future benefits or costs or for the actions of other players. Following Gisser and Sanchez [3] the myopic strategy is obtained by equating marginal benefit to marginal cost for a single player and solving for the pumping rate independently at each stage. The coefficient matrices are described in appendix B.5.

3.3 Results and discussion

The common pool problem formulation outlined above enables us to address several research questions related to the general theme of aquifer depletion. The emphasis here is on issues that cannot be readily considered with simplified bucket models but can be investigated by applying computationally convenient closed form solution methods to the multi-compartment model pictured in Figure 3-1. The questions of particular interest include:

- How does aquifer depletion change over time and how does pumping strategy affect depletion dynamics?
- How much does revenue efficiency decrease with an uncooperative pumping strategy, as compared to a cooperative or social optimal strategy?
- How is revenue efficiency affected by player asymmetry and by the number of players exploiting the groundwater resource?
- How are depletion strategies affected by aquifer stock and river flux limitations?

These questions address interactions between physical constraints described by the groundwater model and management aspects that depend on the benefits and costs experienced by the various players and on their strategies for exploiting the resource. Physical and management factors are explicitly integrated by using a game/optimization approach to describe the groundwater extraction problem.

For specificity, we examine the above questions with a hypothetical example based on a semi-arid irrigated region that is similar in some respects to California's San Joaquin Valley. This example portrays a cross-section of an agricultural valley (overlying the inner aquifer of the compartmental model) that receives minimal rainfall during the growing season and is recharged from a larger non-irrigated upland area (the outer aquifer). Both aquifers are assumed to be unconfined. Some of the inner aquifer groundwater may flow to or from a river running down the center of the valley. The cropland area is divided, for purposes of this example, among players who represent different irrigation districts that compete for the common stock of groundwater.

Parameter	Symbol	Representative value
Outer Aquifer (recharge) area	A_O	7200 km^2
Inner Aquifer (cropland) area	A_I	3600 km^2
Constant lateral flow coefficients	$\lambda_{OI}, \lambda_{IR}$	$32.8, 9.8\ km^2 yr^{-1}$
Hydraulic conductivity	κ_{OI}, κ_{IR}	$100, 100m/day$
Width (aquifer cross-section)	w	$80km$
Traverse length	$\Delta x_{OI}, \Delta x_{IR}$	$80km, 40km$
Specific yield (unconfined)	θ	0.1
Recharge rate	R	$100mmyr^{-1}$
Ground height	h_g	270 m above MSL
River head	h_R	200 m above MSL
Aquifer bottom elevation	h_b	0 m
Single stage time step	Δt	$1yr$
Number of stages	K	60 yr
Discount rate	r	3%
Number of players	N	2
Benefit function parameters	b_1^i, b_2^i	$0.1\$m^{-1}km^{-2}, -3.5 \times 10^{-4}\$m^{-2}km^{-4}$
Pumping Cost parameter	η, c_p	$50\%, 0.12\$kWhr^{-1}$

Table 3.1: Compartmental model input values

We focus on long-term groundwater depletion due to pumping in the inner aquifer. Each player operates many individual wells scattered over a specified crop area. A player's/district's pumpage at any given stage is the total water removed from its wells. The wells are assumed to be distributed with sufficient density to justify the approximation that pumping costs are determined by the spatially and annually averaged inner aquifer head (see discussion in Section 3.2.1 above). In our nominal case we assume that the lateral flow coefficients are constant, which implies that the aquifer depth is large compared to head drawdown. The nominal benefit function coefficients reflect the derived demand for pumped groundwater for growing a water-intensive crop such as almonds, grown by all players. Table 3.1 summarizes the nominal values for the compartmental model inputs similar to the San Joaquin valley [58]. Some of these are modified in subsequent sensitivity analyses.

3.3.1 Dynamical behavior - Nominal case with two identical players

In the nominal case we consider a decline in water level and storage from a natural unpumped steady state to a new pumped steady state when two identical players compete for groundwater with constant lateral flow coefficients between compartments. Figure 3-2 shows aquifer water level and total (sum over all players) pumping time histories for the four decision strategies summarized in Section 3.2.2. Note that

it takes a few hundred years for the aquifer system to reach a pumped steady state in the nominal case. Figure 3-3 plots the results over a shorter 60 year time scale that gives a better picture of pumping and storage dynamics. The pumping rate starts at a high level, when pumping costs are low, and gradually decreases as the inner aquifer water level falls and costs increase. The same basic pattern applies for all four decision strategies, with the social optimum giving somewhat less pumping and drawdown and the myopic strategy giving the most depletion.

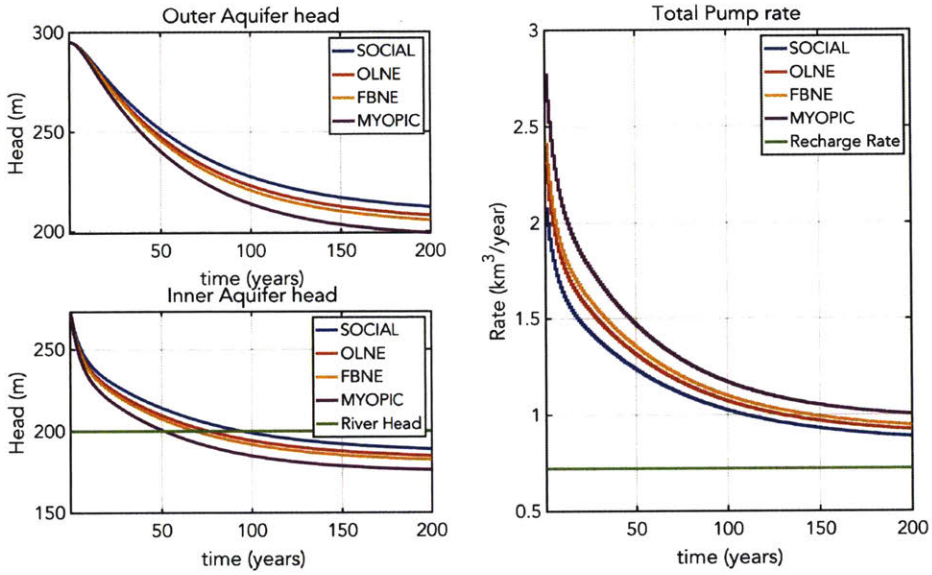

Figure 3-2: Time series simulation of the lumped two-state groundwater system for the nominal case over a 200 year timescale. From an unpumped steady state, the water level reduces to a lower steady state indicated in the top and bottom plots of left panel which show outer and inner aquifer respectively. The right panel shows the total pump rate with different strategies in comparison to the total recharge obtained from the outer aquifer.

The decline in pumping rate shown in this example is typical of optimization and game analyses based on decreasing marginal demand and increasing marginal cost for pumped groundwater. Players can be expected to adjust to this decrease in water availability in a number of ways. They can reduce the fraction of cropland that is irrigated, switch to more water efficient and/or more profitable crops or improve irrigation efficiency. Another possibility in some situations is to expand pumping into new previously unpumped areas that are sufficiently far from existing wells for there to be little additional drawdown in the existing pumped area. Either existing or new players could farm the new areas. In any case, the combination of decreased pumping in the original well area and expansion of pumping into new areas can lead to a temporary increase in total pumpage, until there is no further possibility for expansion and pumping rates in the new wells decrease from their initial high values. This expansion option is not considered in the problem formulation given here since

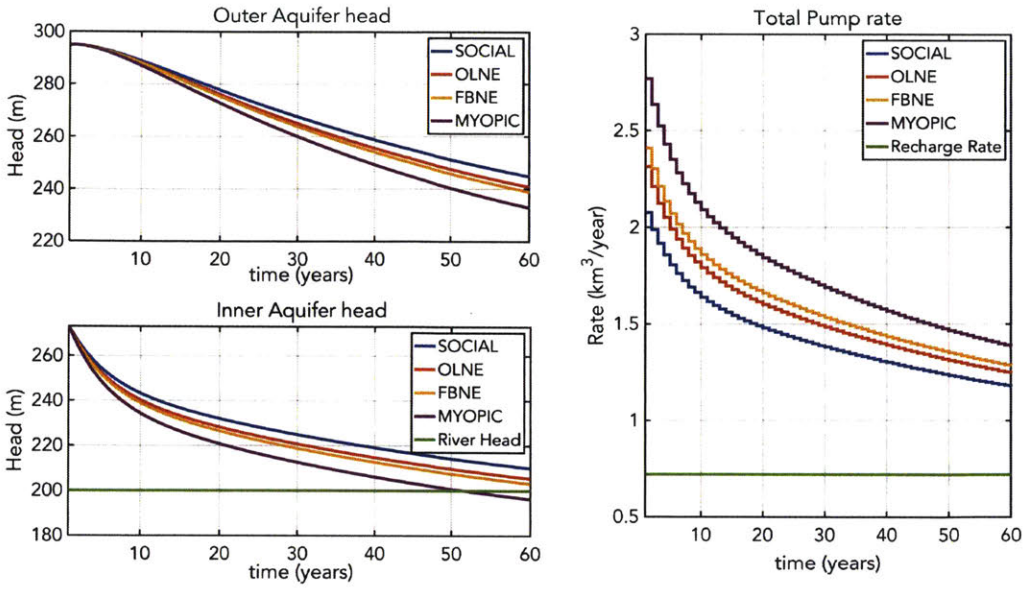

Figure 3-3: Closeup view of the first 60 years of simulation in the nominal case (figure 3-3) to emphasize the transient dynamics

the use of an average compartmental head for cost computations assumes that the well density is already high. However, extensions that provide for expansion of pumping into new areas could be important in some applications.

The multi-compartment model accounts for the fact that the players can draw water from both the recharge area and the river [13]. In our nominal case the decreasing inner aquifer head eventually falls below the river head indicated by the horizontal green line, for all decision strategies. At that transition point the groundwater direction at the river changes from out of the aquifer to into the aquifer. The pumping rate is initially much higher than recharge because water is being withdrawn from storage. It gradually approaches a steady-state value above the recharge value because water is also being drawn from the river. Flow from the river accounts for about 20% of the total steady-state pumpage for the nominal case across the decision strategies.

3.3.2 Revenue efficiency for alternative decision strategies

The inefficiency incurred by non-cooperative and myopic strategies can be measured in terms of the revenue lost, as compared to a socially optimum strategy that maximizes total revenue for a given resource. Another useful performance measure is the increase in steady-state groundwater pumpage above the social optimum value. The first row of Table 3.2 summarizes these indicators for the four different decision strategies for the nominal case. The percentage values in the table indicate revenue efficiency decreases and pumping rate increases for each option. The economic inefficiencies incurred by non-cooperative management of the reservoir in our examples are greater than reported by Gisser and Sanchez but are still within the range of results cited

Case	Revenue (NPV) ($\times 10^9$\$)				Steady state pump rate ($km^3 yr^{-1}$)			
	Social	OLNE	FBNE	Myopic	Social	OLNE	FBNE	Myopic
Nominal	2.3	2.259 1.8%	2.22 3.2%	2.073 9.9%	0.8552	0.8925 4.4%	0.9148 7%	0.9735 14%
Variable recharge	2.19	2.14 2.8%	2.13 3%	1.98 9.7%	-	- -	- -	- -
Asymmetric players	2.76	2.71 2%	2.68 2.9%	2.52 8.7%	0.8894	0.9517 7%	1.006 13%	1.078 21%
Multiple (8) players	2.3	2.155 6.3%	2.135 7.2%	2.073 9.9%	0.8552	0.9165 7.7%	0.9468 10.7%	0.9735 14%

Table 3.2: Performance assessment of non-optimal management of the groundwater system with respect to the social optimum. Percent values indicate relative difference to the social optimum.

by Koundouri [40]. The maximum inefficiencies for the Nash feedback and myopic solutions are 7.2% and 9.9% , respectively, for the 8 player case.

The Nash equilibrium and myopic strategies are more aggressive than the social optimum in drawing down the groundwater stock during the most dynamic period. The steady state heads and pumping rates shown in Figure 3-2 maintain the relative ranking of decision strategies that developed during the transient period. The importance of short-term behavior in establishing the degree of groundwater depletion below natural conditions emphasizes the need to look at dynamic as well as steady state solutions. Also, the early time phase with rapidly changing heads and pumping rates is likely to be of more practical interest since the climatic and economic factors that favor particular initial pumping rates could easily change over the long period needed to reach steady state. The time to steady state depends significantly on physical properties such as the ratio between the recharge and pumping areas and the coefficients that control lateral flow. These properties are not included in the bucket model.

The open loop Nash solution pumps less and gives less depletion than the feedback or myopic solutions since all players commit to a pumping schedule in advance, without deviation. This commitment reduces player flexibility but brings the open loop revenue efficiency closer to the social optimum than these other options [64]. On the other hand, the open loop strategy lacks a real-time ability to respond to unforeseen fluctuations in inputs such as recharge. This is apparent in Figure 3-4, which gives head and total pumping trajectories obtained for a hypothetical changing recharge history, shown in green on the pumping plot. The social optimum, open loop, and feedback strategies are derived assuming that recharge remains constant at the nominal value. Since the social optimum and feedback, as well as the myopic strategy, obtain releases from a real-time decision rule that depends on the current state they are able to adjust their pumping rates when recharges different than those assumed in the decision rule derivation. By contrast, the open-loop trajectory remains the same as if the recharge were constant and, as a result, tends to overpump in dry periods and

underpump in wet periods. The result is that the open loop Nash equilibrium strategy revenue performance reduces to the level of feedback Nash equilibrium strategy (see Row 2 of Table 3.2).

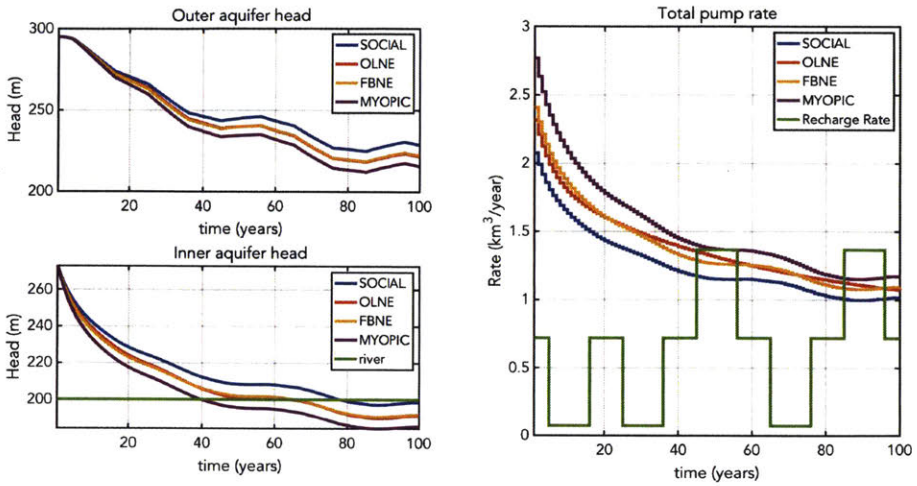

Figure 3-4: Time series simulation of the heads and the pump rates under variable recharge conditions. Social optimum, Feedback Nash Equilibrium and Myopic strategies dynamically adapt their decisions to unpredictable water level scenarios. Open loop Nash equilibrium sticks to the pre-committed strategy and performs worse.

3.3.3 Asymmetrical benefit functions and multiple players

Our nominal case is based on two players who generate about 50000 $\$km^{-2}yr^{-1}$ using about $800mm$ of water per year ([34]). But players often grow different crops with different yields, market prices, and water requirements. As an example we consider how pumping is affected when one of the players grows a crop that can generates 20% more revenue per unit area than the crop considered in the nominal case with the water requirement unchanged. The other player grows the nominal crop. Figure 3-5 shows on the left the resulting inner and outer aquifer heads and on the right the individual pumping rates of the two players. The ordering of decision strategies remains the same in terms total pumpage and the transient head behavior.

An important feature to notice in the individual pumping histories is the relative distance of the non-cooperative strategies (OLNE and FBNE) from the myopic and the social optimum strategies. The reduced impact of the player with the low value crop tends to make this player adopt a more short-sighted (nearly myopic) strategy.

A number of studies have suggested that the revenue inefficiency associated with uncooperative groundwater management increases as more players compete for the resource [63]. A two player formulation, while analytically convenient, is often unrealistic. Fortunately, it is straightforward to solve the nominal problem formulated in Section 3.2 for a reasonably large number of players if the constraints have constant lateral flow coefficients and are affine. Here we consider this case for from 1

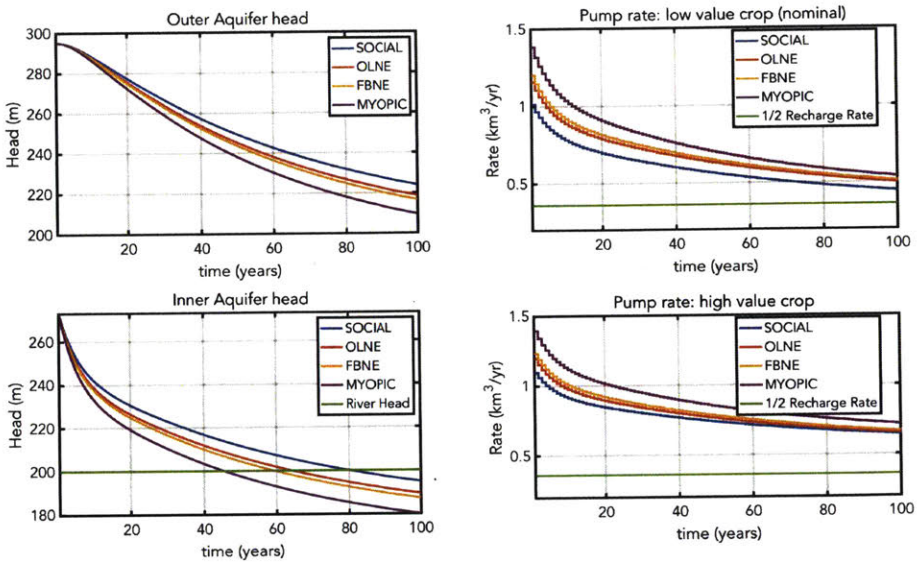

Figure 3-5: Head and total pumping time series for two players with asymmetric objective functions. Top right panel gives pumping histories for the player with the lower value crop while lower right panel gives the histories for the player with the higher value crop. Different strategies pump different quantities of water to maximize their net present revenue.

to 8 identical players, using the nominal inputs from Table 3.1 and a feedback Nash equilibrium decision strategy (which is the same as the social optimum for $N = 1$). The results verify that as the number of players increases the inner and outer aquifer heads drop and pumping increases. As the number of players increases the future impact of an individual player's action on the system becomes insignificant, leading to a nearly myopic strategy. At $N = 8$ the feedback Nash and myopic solutions are barely distinguishable. Similar results are obtained for the open loop strategy (not shown). This comparison suggests that the uncooperative feedback revenue efficiency decreases as the number of players grows but is bounded from below by the myopic value.

3.3.4 Stock and river flux limitations

The variations on the nominal example considered in the above sections all assume that pumping costs are sufficiently high to limit drawdown to a small fraction of the aquifer depth, which is defined as the difference between the groundwater head and the aquifer bottom elevation. In this case the unconfined lateral flow coefficients can be approximated by constants. When pumping costs are low, pumping rates are high, and the aquifer depth is relatively small, these coefficients should decrease to reflect the fact that higher head gradients are needed to supply a given flux of water ([70],[9]). An approximate way to handle this is to assume the flow coefficients are

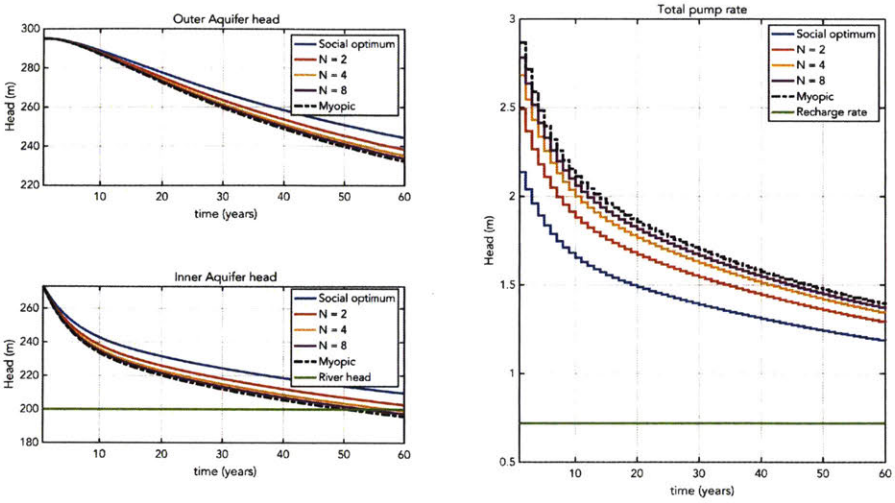

Figure 3-6: Time series simulation of closed loop Nash equilibrium solution when the number of agents increases (N=1,2,4,8). As the agents increase the inefficient groundwater pumping increases till it reaches the asymptotic pumprate of the myopic strategy

time-dependent and linearly proportional to the average flow depth of the associated compartments:

$$\lambda_{OI,k} = \frac{\kappa_{OI} w}{\Delta x_{OI}} \left(\frac{h_{I,k} + h_{O,k}}{2} - h_b \right); \quad \lambda_{IR,k} = \frac{\kappa_{IR} w}{\Delta x_{IR}} \left(\frac{h_{I,k} + h_{R,k}}{2} - h_b \right) \tag{3.14}$$

where h_b is the aquifer bottom elevation (common to both aquifers), w is the width of the aquifer cross section, Δx_{OI} and Δx_{IR} are distances between the centers of the outer and inner aquifer and the inner aquifer and the river, and κ_{OI} and κ_{IR} are effective hydraulic conductivities across the two interfaces. In this formulation the flow depths and lateral flow coefficients decrease as the head drops in response to pumping. Additional inequality constraints insure that the heads remain above the aquifer bottom elevation [59]:

$$h_{I,k} \geq h_b \quad h_{O,k} \geq h_b \tag{3.15}$$

Another relevant constraint is the limited capacity of the river to act as a water source when the inner aquifer head is drawn down. In particular, we suppose in this example that the flow $Q_{IR} = \lambda_{IR,k}(h_{Ik} - h_{Rk})$ withdrawn from the river as a result of pumping cannot exceed a specified fraction f of the upstream inflow:

$$Q_{RI} \leq f Q_{Rinf} \tag{3.16}$$

The exact fraction of upstream inflow allowed could vary based on regulatory conditions but cannot exceed $f = 1.0$.

If the lateral flow coefficients of 3.14 are used instead of the constants used in the nominal case and the inequality constraints of 3.7,3.15, and 3.16 are added it is no longer possible to use the closed form Ricatti-based solutions given in the Appendix to find the open loop and feedback decision strategies. Numerical solutions can be derived using an iterative dynamic programming approach but existence and uniqueness of the Nash equilibria are more difficult to demonstrate.

Fortunately, it is relatively straightforward to derive the single player (single objective) social optimum and myopic solutions when 3.7,3.14,3.15, and 3.16 are included if we use a standard interior point algorithm to solve the resulting nonlinear optimization problem. This more limited analysis (which omits the open loop and feedback Nash strategies) can still be quite informative since the social optimal and myopic head and revenue solutions are typically upper and lower bounds, respectively, for the corresponding game solutions. The ordering of social optimum-open loop-feedback-myopic solutions (from highest to lowest head and revenue) is apparent in all of the results described above. A similar ordering is described by Provencher and Burt [59] and Rubio and Casino [63] in their steady-state analyses of groundwater depletion. Given these observations, we examine here the results obtained for the social optimum and myopic strategies when the flow coefficients have the head-dependent form given in 3.14, the state equations are quadratic rather than affine, and the inequalities 3.7, 3.15, and 3.16 are added. These extensions provide an useful insight about the role of stock and river flux limitations in the common pool problem.

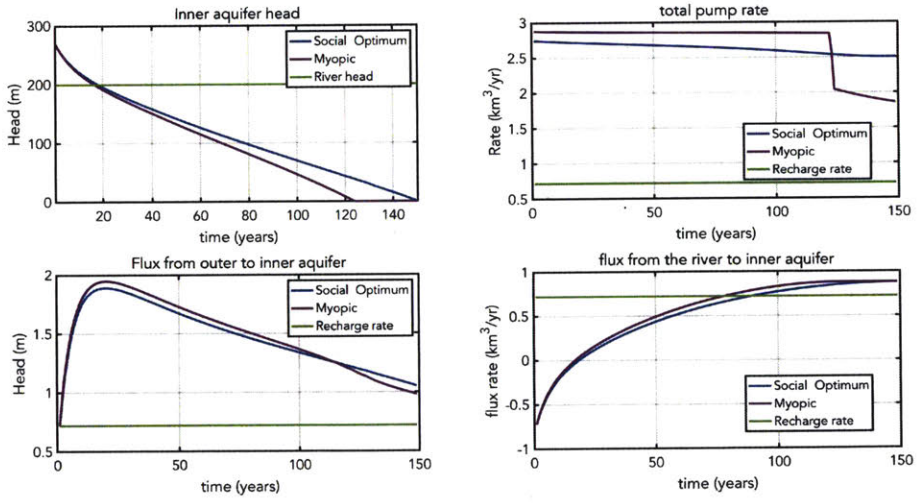

Figure 3-7: Time histories of the water level in the outer and inner aquifer and pump rate with a non-linear lateral flow coefficient model and finite stock-flux constraint. With a very low cost for pumping, the social and myopic strategy choose to drain the inner aquifer head until they reach its bottom. At the aquifer bottom both strategies are limited to maximum flux from the river equal to $flux_{max} = \frac{\kappa_{IR} w}{\Delta x_{IR}} \left(\frac{h_R}{2} - h_b\right) h_R$

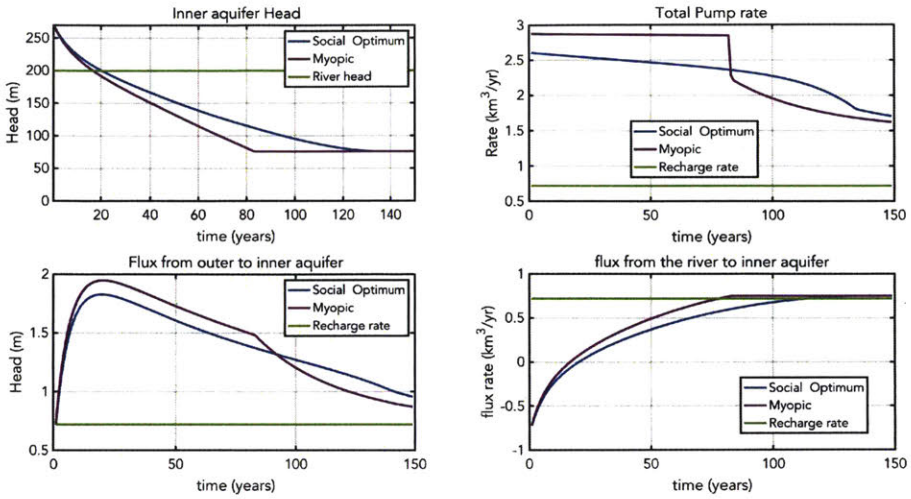

Figure 3-8: Time histories of the water level in the outer and inner aquifer and pump rate with a non-linear lateral flow coefficient model and finite stock-flux constraint. The social and myopic strategy do not reach the aquifer bottom, but hit the maximum river-to-inner aquifer flux constraint $flux_{max} = fQ_{Rinf}$

To illustrate the significance of the stock and flux limitations, we use the values for $h_b, \kappa_{OI}, \kappa_{IR}$ given in Table 3.1 and reduce the pumping cost to $0.001\$kWhr^{-1}$, while keeping all other inputs at their nominal values. We consider two cases: the first with an upstream river inflow that is much greater than the total maximum irrigation demand and the second with an upstream inflow that is comparable to this demand. Figure 3-7 shows the time series simulation of a high upstream flow case ($f = 0.1$, $Q_{Rinf} = 20km^3yr^{-1}$). The lowered pumping cost allows the strategies to pump higher volumes until they reach the aquifer bottom [$h_I = h_b = 0$] (See top panel figure 3-7). Beyond this point the players can no longer increase the flux from the river to the inner aquifer due to their fixed head difference. The inner aquifer stock constraint of 3.15 becomes active and the river flux constraint of 3.16 remains inactive. The total pumpage will eventually reach a steady state equal to the sum of the total recharge and the fixed nominal flux from the river given by;

$$u_{SS} = RA_O + \frac{\kappa_{IR} w}{\Delta x_{IR}} \left(\frac{h_R - h_b}{2} \right) h_R \tag{3.17}$$

Figure 3-8 shows the second low river inflow case with an upstream inflow rate pf $Q_{Rinf} = 0.75km^3yr^{-1}$). This corresponds to $f = 1$. The river flux constraint now becomes active and the stock constraints remain inactive (see left top panel). Both strategies eventually reach a steady state pumping rate equal to the sum of the recharge rate and the flux constraint given by:

$$u_{SS} = RA_O + fQ_{Rinf} \tag{3.18}$$

These physical limitations are not revealed in an analysis that only uses cost to limit pumpage.

3.4 Conclusions

This chapter describes groundwater depletion as a multi-player non-cooperative dynamic game. The game is solved for various decision strategies with appropriate optimization methods. Open loop and feedback Nash uncooperative game solutions are solved with a convenient closed form approach that applies for affine-quadratic player objective functions and affine equality constraints. The constraints are obtained from a multi-compartment groundwater model formulated as a set of affine state equations. This model, which represents the common pool resource, captures essential physical features that affect the game solution. They include the existence of a credible unpumped hydrologic steady-state, provision for realistic boundary conditions, and stock and flux restrictions on depletion.

The example discussed in Section 3.3.1 provides an informative comparisons of social optimum, open loop Nash, feedback Nash, and myopic decision strategies. Comparisons are made for different problem formulations that consider asymmetric player benefits, multiple players, and variable recharge. A single player following the social optimum decision strategy consistently generates the largest net present value revenue from the resource over a range of different situations. The social optimum is followed by the open loop Nash equilibrium and a feedback Nash equilibrium, with the myopic strategy giving the least revenue. However, the differences between these strategies are sometimes small and never larger than 10% of the social optimum revenue for the example considered here.

The open loop strategy does somewhat better than the other non-cooperative strategies because it adheres to a pre-committed pumping trajectory that enables players to take actions that may be detrimental in the short term in order to maximize long term benefit. The feedback solution does not benefit from a prior commitment but is better able to adapt to changing conditions such as variable recharge or changes in crop value and energy costs. This is because it derives pumping rates in real-time based on currently observed hydraulic heads. The open loop and feedback Nash strategies give results closer to the short-term myopic strategy when players have asymmetric revenue functions or when a large number of players are sharing the same aquifer.

Our analysis of the two Nash equilibrium strategies relies on closed form game solutions that require all constraints to be affine equalities. This requirement can be met to an acceptable degree of accuracy if pumping costs and aquifer geometry maintain a large and reasonably constant aquifer flow depth (the difference between the hydraulic head and the base elevation). When the inner aquifer depth decreases significantly from its unpumped value the lateral flow coefficients in the compartmental equations can no longer be assumed to be constants and stock and flux limitations expressed as inequality constraints may become active. In this case the affine state equations need to be replaced by a combination of nonlinear equalities and linear in-

equalities. The associated closed form Nash solution approach no longer applies and the Nash equilibrium solutions must be found iteratively. Convergence and uniqueness of these solutions can be difficult to prove. In our analysis of the low aquifer flow depth case we consider only the social and myopic strategies, which bracket the open and feedback Nash equilibrium strategies in terms of efficiency and pumping rate. An analysis of these strategies for our example problem indicates that stock and flux limitations can play an important role in determining the maximum water available for pumping when the unit pump cost is too low to be limiting.

Taken together, these investigations provide valuable insight about player and groundwater flow behavior over a range of decision strategies, hydrologic conditions, and economic parameters. The optimization results presented here can support water management policies by indicating when economic inefficiencies due to uncooperative exploitation are significant and when they can probably be neglected. Of course, the model and example considered here are still quite simplified compared to reality and caution should be used in extrapolating our results, especially in support of policy recommendations. A game analysis is most useful for framing the problem, suggesting hypotheses, and indicating areas of uncertainty that require additional attention. In these respects it provides a good way to gain better understanding of the factors responsible for undesirable groundwater depletion.

Chapter 4

Managing groundwater as a common pool resource: A spatially distributed approach

4.1 Introduction

Chapter 3 discusses the consequences of groundwater pumping under different management strategies using a spatially lumped aquifer model. This analysis improves on studies that rely on a bucket model by using a multi-compartmental approach that provides a more realistic treatment of boundary conditions, flow in the absence of pumping, and the effects of stock and river flux limitations. The approach provides useful quantitative estimates of inefficiencies due to uncoooperative aquifer exploitation. However, there are certain drawbacks to a spatially lumped analysis. In particular, in a lumped model the players in a given compartment share the same spatially averaged head, which depends only on the total pumping rate summed over all players. In reality, heads vary significantly over space and the head at each player's well might depend primarily on that player's own pumping rate and only secondarily on the pumping rates of others. The influence of other players' pumping rates on a given player's well drawdown and pumping cost depends strongly on the geometric arrangement and proximity of the wells. When wells are sufficiently far apart the impact of uncooperative management is small. When they are closer together the impact can be greater ([14]). These important spatial effects are not considered in a lumped groundwater model but can be included in a higher resolution spatially distributed model.

Recent literature on the groundwater common pool problem has placed increased emphasis on using more realistic physical models and/or descriptions of player behavior ([56],[50],[70],[57],[31]). Some of this research tends to emphasize economic aspects and uses simplified analytical solutions to describe groundwater flow. These studies can account for the strategic externalities and economic inefficiencies that arise during uncooperative exploitation of the resource but they generally cannot provide physically convincing descriptions of groundwater flow ([3],[63],[56],[57]). Other

research studies emphasize hydrologic aspects and use more realistic groundwater models but are limited in their ability to account for the behavior of uncooperative players who anticipate future consequences and respond to other players' actions in real time. These studies typically focus on social optimum or myopic strategies or rely on multi-agent system simulation (MASS) ([14],[50],[70]).A few studies combine realistic groundwater solutions based on superposition of individual well drawdown solutions or aquifer response functions with optimization or game theoretic methods suitable for particular decision strategies or problem formulations ([14],[46],[49]). Our objective in this chapter is to further bridge the gap between economically-focused and hydrologically-focused approaches by incorporating a more general high resolution spatially distributed groundwater model into a convenient closed form approach for solving uncooperative games.

The spatially distributed model used here is based on a space and time-discretized (finite element) partial differential equation (PDE) description of unconfined groundwater flow. This PDE system can accommodate heterogeneity in hydrogeological parameters such as transmissivities and boundary conditions to provide a tool that is able to represent the distinctive characteristics of particular aquifers and management problems. We use the spatially distributed model to compare the performance of three of the pumping strategies already considered in Chapter 3: Social optimum, Feedback Nash Equilibrium, and myopic. The comparison is based on a simple hypothetical example that clarifies the important role of the player well configuration and of connections between groundwater sources and well locations.

4.2 Formulation of the spatially distributed groundwater model

4.2.1 Objective and constraints

We formulate the spatially distributed common pool problem as an N-person finite horizon non-cooperative discrete time deterministic game, following the notation of Basar and Oldser [8]. The objective functional to be maximized by Player i is:

$$\begin{aligned} L^i(u^1,...,u^N) &= \sum_{k=1}^{K-1} g_k^i(u_k^1,...,u_k^N,x_k,x_{k+1}) + g_K(x_K); \quad i=1,...,N \\ &= \sum_{k=1}^{K-1}(1+r)^{-k}\left[B_k^i(u_k^i) - C_k^i(u_k^i,x_k)\right] + g_K(x_K) \end{aligned} \tag{4.1}$$

where u_k^i is the control variable (water pumped, L^3T^{-1}) for Player i over stage (time step) k and u^i is a vector of Player i's pumping rates for all K stages. For convenience we suppose that each player operates a single well at a specified location. This assumption can be readily generalized to allow for multiple wells for each player, if desired. The discounted net benefit in each stage g_k^i (\$) is proportional to the difference between the current benefit $B_k^i(u_k^i)$ from using the pumped groundwater to

grow crops and the current cost of pumping $C_k^i(u_k^i, x_k)$, with a discount rate r. In the spatially distributed formulation the pumping cost for a given player is proportional to the local drawdown at that player's well location. This is in contrast to the lumped formulation of Chapter 3, which computes player cost from a spatially averaged compartmental head that is common to all players.

The state equation is derived from the following vertically-averaged unconfined aquifer equation [70]:

$$\theta \frac{\partial h(x,t)}{\partial t} = \nabla.(\kappa(h(x,t) - h_B)\nabla h(x,t)) + q(x,t) \tag{4.2}$$

where κ is the aquifer hydraulic conductivity, h_B is the elevation at the base of the aquifer, $q(x,t)$ is a forcing function that defines the pumping withdrawal distribution over space and time, and θ is the aquifer specific yield. This equation can be assumed to be linear when the unconfined aquifer depth $h(x,t) - h_B$ is sufficiently large to approximate $\kappa(h(x,t) - h_B)$ by a head-independent transmissivity $T(x,t)$. We restrict our attention to situations where this approximation is reasonable since it enables us to use a convenient closed form solution for feedback Nash equilibrium management strategy.

The initial head is specified over the domain Ω at $t = 0$:

$$h = h_0 \; in \; \Omega \; t = 0 \tag{4.3}$$

A mixed boundary condition is specified over the domain boundary Υ as follows:

$$\kappa \frac{\partial h}{\partial n} = \alpha(H(x) - h(x,t)) + Q \tag{4.4}$$

Where H is a prescribed time invariant head condition at specified locations along the boundary such as a river or lake and Q is a specified constant flux.

We suppose here that the aquifer transmissivity is spatially uniform so the groundwater flow state equation can be reduced to:

$$\theta \frac{\partial h(x,t)}{\partial t} = T\nabla^2 h(x,t)) + q(x,t) \tag{4.5}$$

The transmissivity can be allowed to vary over space, if desired, without affecting the basic problem formulation. Equation 4.5 can be solved using numerical techniques that rely on a discrete spatial discretization grid such as the one illustratated in Figure 4-1. The finite element grid shown in the figure approximates the continuously varying hydraulic head with a piecewise linear function that interpolates the head values at a finite number of node points at the corners of the discrete grid cells (elements). These nodal head values are the unknowns in the discretized finite element solution procedure. If time is also discretized to be compatible with 4.1 the partial differential equation is approximated by the following set of discrete difference equations:

$$h_{k+1} = F_k h_k + G^1 u_k^1 + ... + G^N u_k^N + E_k \tag{4.6}$$

where the matrices F_k, G_k^i are obtained from the time and space discretization of the

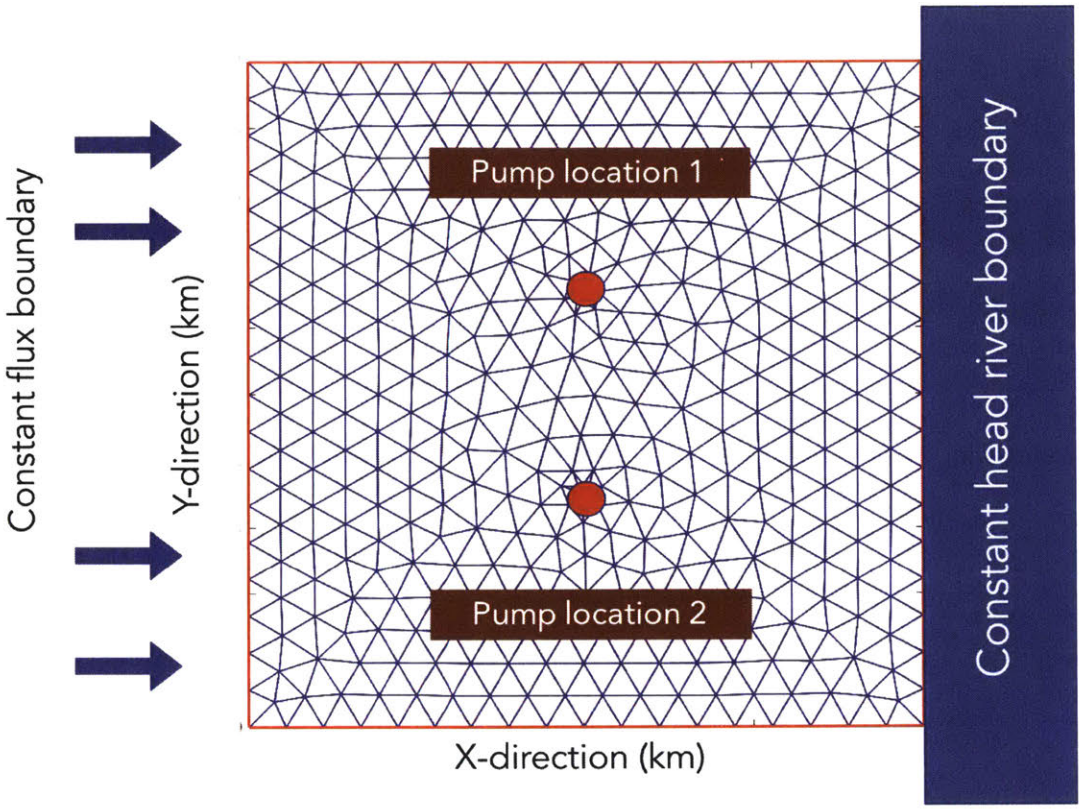

Figure 4-1: Mesh discretization of a spatially distributed groundwater showing the pumping locations that act as sink terms and the boundary conditions

PDE system. Further details about generating the matrices are provided in appendix C.

The player benefit functions used in the spatially distributed problem formulation have the same quadratic structure used in the lumped model of Chapter 3:

$$B_k^i(u_k^i) = b_1^i u_k^i + b_2^i (u_k^i)^2 \tag{4.7}$$

$$0 < u_k^i < u_{max}^i = -\frac{b_1^i}{2b_2^i} \; \forall \; i \in N \tag{4.8}$$

However, the pumping cost functions are different from the lumped model versions. The cost for Player i depends on the depth to groundwater at that player's particular well location. This can be described compactly as follows:

$$C_k^i(u_k^i, h_k) = c\eta(h_g - m_i^T h_k)u_k^i; \; c > 0 \tag{4.9}$$

where h_g is the specified ground surface elevation. The vector m_i has a few non-zero elements and the rest equal to 0. This selects out the component of the discretized head vector at player i's well location. For simplicity we assume that the ground surface elevation is constant (i.e. the same at all wells). With this formulation each

player's pumping cost depends on the local drawdown at his/her pumping well. The drawdown is no longer the same at all the pumping locations or for all players.

The objectives and state equation given above can be put in an affine quadratic form that permits us to use closed form solutions for the social optimum, feedback Nash equlibrium, and myopic decision strategies. All three of these strategies give affine decision rules that depend on the current state of the system. The affine quadratic objective function can be expressed as:

$$L^i(u^1, ..., u^N) = {h_K}^T Q^i_{hh,K} h_K + \sum_{k=1}^{K-1} {u^i_k}^T Q^i_{uu,k} u^i_k + {h_k}^T Q^i_{uh,k} u^i_k + S^i_{uk} u^i_k \quad (4.10)$$

$$h_{k+1} = F_k h_k + G^1 u^1_k + ... + G^N u^N_k + E_k \quad (4.11)$$

and the state-dependent affine decision rule is:

$$u^i_k = \gamma^i_k(x_k) = \Gamma^i_k h_k + \Upsilon^i_k \quad (4.12)$$

Here the pumping decision is dependent not only on the local well head but also on the heads at all the other nodes in the aquifer, including those pumped by other players.

4.3 Results and discussion

With a spatially distributed formulation of the groundwater system, the questions of strategic externality and non-cooperation examined in Chapter 3 can be re-visited. In particular, in this section we address the following research questions:

- How does the aquifer head vary over time and space and what is the influence of well geometry on the head distribution ?
- How much does the revenue efficiency decrease with an uncooperative pumping strategy, as compared to a cooperative or social optimal strategy?
- How do the fluxes between water sources (boundaries) and sinks (wells) vary over time and space and how do they depend on decision strategy and well geometry?

Analysis of these topics depends on the specific properties and capabilities of the physical systems model and on the objectives and decision strategies of the players.

As a nominal case, we examine the above questions in a small rectangular domain ($400km^2$) to observe local head drawdowns and fluxes. Figure 4-1 shows this domain, which contains two pumping wells belonging to Players 1 and 2. It has a zero flux boundary condition on the top and bottom edge boundaries. On the left boundary is a constant prescribed flux Q. On the right boundary is a river modeled as a constant head condition with spsecified head h_R. Each of two identical players cultivates crops on his/her half the of the domain. Table 4.1 summarizes the nominal values used for the parameters.

Parameter	Symbol	Representative value
Domain Area	d_{area}	$400km^2$
Domain length	d_l	$20km$
Domain width	d_w	$20km$
Hydraulic conductivity	K	$1km/yr$
Specific yield (unconfined)	θ	0.1
Prescribed flux	Q	$0.2mkm^{-1}yr^{-}1$
River head	h_r	$0m$ above datum
Single stage time step	Δt	$1\ yr$
Number of state	K	$50\ yr$
Discount rate	r	3%
Number of players	N	2
Benefit function parameters	b_1^i, b_2^i	$0.1\$m^{-1}km^{-2}$, $3.2\$m^{-2}km^{-4}$
Pumping cost parameter	η, c	$50\%, 0.15\$kWhr^{-1}$

Table 4.1: Representative values for the parameters

4.3.1 The effect of well spatial configuration

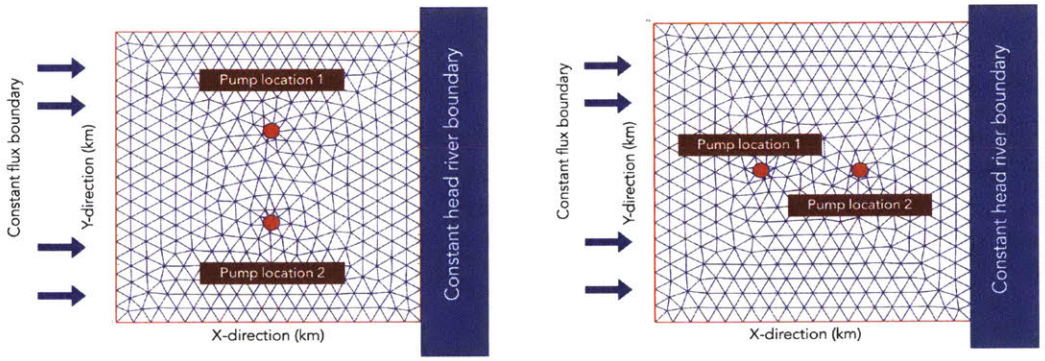

Figure 4-2: Flow and head behavior with two different well configurations. Left shows a symmetric placement with respect to the boundaries while the right figure shows an asymmetric placement

In this section we compare the flow behavior between the two well configurations shown in Figure 4-2 . The left panel shows an symmetric configuration where the pumps are located equidistant from the two boundaries. The two wells have equal access to the constant recharge and the constant river head boundary. The right panel shows an asymmetric well configuration with one well closer to the recharge and the other closer to the river head boundary. In the absence of pumping wells, the prescribed recharge flux for both configurations flows from the left to the right boundary. Figure 4-3 shows the corresponding head contours and flow lines.

When the wells are pumped the groundwater system eventually reaches a new steady state with the head generally lower than the initial unpumped solution. This new steady-state depends on the decision strategy used to determine the pumping

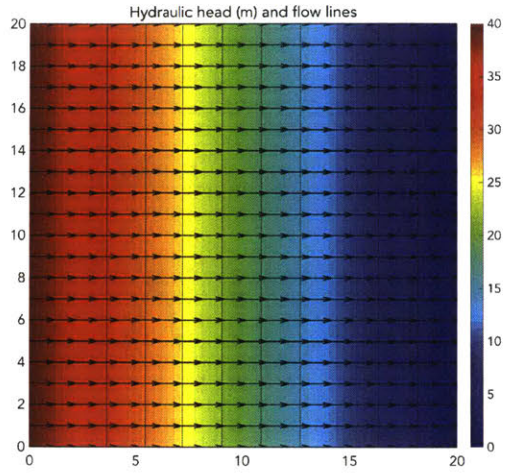

Figure 4-3: Hydraulic head condition and flow lines in the absence of pumping wells

rates. The steady-state groundwater flow distributions for the feedback Nash equilibrium strategy are discussed below in a separate section for each well configuration.

Symmetric Configuration Results

Figure 4-4 shows the steady state head contours and flow directions for the symmetric well configuration when pumping is based on an uncooperative feedback Nash strategy. The head lines reveal local drawdown at the two pump locations near the top (player 1) and bottom (player 2). The head gradients drive the water from the flux boundaries into the wells. The pumps draw water from the specified head river boundary on the right as well as from the specified flux boundary on the left. The head contour curves and the flow lines are symmetrical since the players share the same marginal benefit and cost functions and the well configuration is symmetrical.

Figure 4-5 shows the steady state pumped head profiles for the three candidate decision strategies along a bottom to top (Y axis) section passing through both player wells. Note that the the pumped steady-state heads are significantly lower than the uniform unpumped steady-state head at this central location. Figure 4-6 shows the head and pumping time histories for the three management strategies. For this completely symmetrical case the two players have identical head and pumping histories and identical revenues and costs. With a local drawdown lower than the river head, it is possible to sustain a steady-state pumping rate higher than the total recharge.

Asymmetric Configuration

Figure 4-7 shows the pumped steady state for an uncooperative feedback Nash equilibrium strategy with the asymmetric well configuration. It is interesting to note that in this case some of the groundwater entering from the left boundary flows past the wells to the right boundary. The well closer to the left boundary (Player 1) obtains

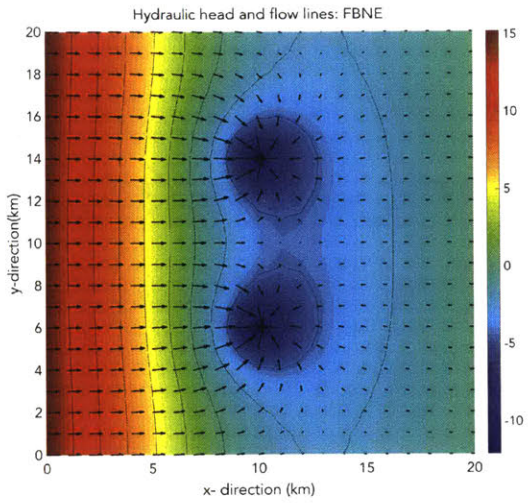

Figure 4-4: Contour plot and flow lines for a symmetric pump geometry using a feedback Nash equilibrium strategy

all of its water from the recharge. The well on the right (Player 2) draws in water from both the recharge and the river head boundaries.

For the nominal pumping cost, the Player 1 pumping rate cannot exceed the total flux available from the left boundary. To obtain additional water from the river Player 1 needs to generate a much higher local drawdown than the drawdown at the Player 2 well. This is an expensive strategy that would drive up the Player 1 pumping cost.

Figures 4-8 and 4-9 illustrates this point with steady-state head profiles along a left to right (x axis) section through the two wells as well as plots of the total pumping history (sum for the two players). At the pumped steady-state, the social optimum and feedback Nash equilibrium strategy give a higher head in the left well than the right. By contrast, the myopic strategy gives a somewhat greater drawdown at Well 1. The pumping rate starts from a high value with lower pumping costs and gradually decreases as the local drawdown increases. This is observed in all the three decision rules, with the social optimum giving the lowest steady-state pumping and drawdown and the myopic giving the highest. Figure 4-8 shows that the total pumping rate is greater than the recharge, reflecting the contribution from the river.

Figure 4-9 distinguishes the pumping rates between the two players. The well closer access to the river boundary pumps more than the total recharge ($Qd_w = 0.04km^3yr^{-1}$) and obtains most of its water from the river. The player closer to the flux boundary is restricted to a lower steady state pump rate due to lack of river access.

4.3.2 Revenue efficiency for alternative decision strategies

The inefficiency incurred by the feedback Nash equilibrium and myopic strategies can be measured in terms of the percentage of revenue lost, as compared to the social

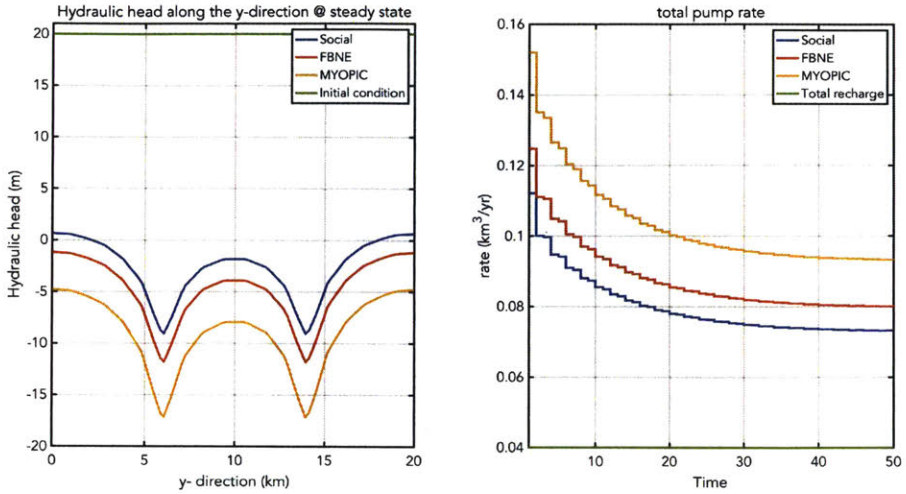

Figure 4-5: Left panel compares the aquifer head profile at steady state along the y direction at $x = 10km$ between social optimum, feedback Nash equilibrium and myopic strategies. Right panel shows the their total pumping histories.

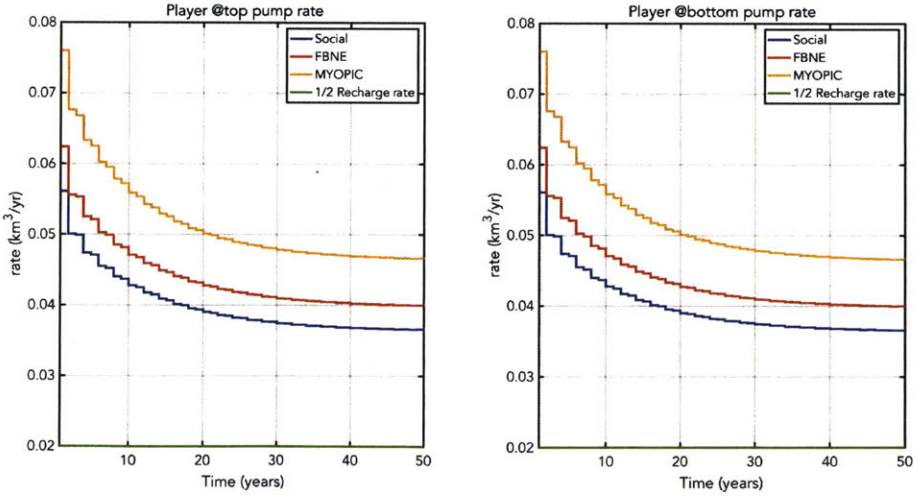

Figure 4-6: Individual player pump rates with asymmetric well configuration under three decision strategies. With equal net benefit functions, players show the same pumping behavior

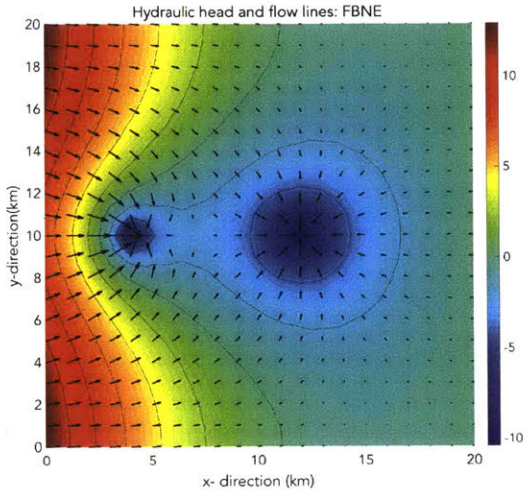

Figure 4-7: Contour plot and flow lines for the asymmetric pump geometry using a feedback Nash equilibrium strategy

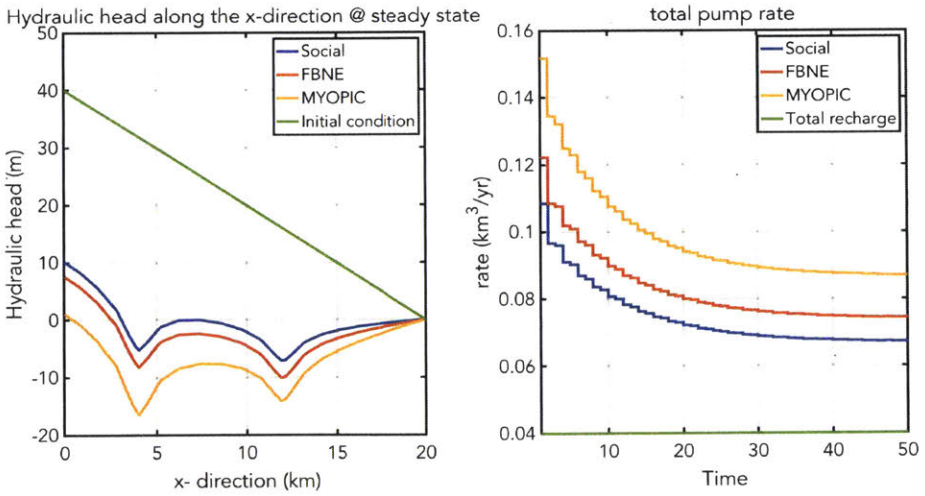

Figure 4-8: Left panel compares the aquifer head profile at steady state along the x direction at $y = 10km$ between social optimum, feedback Nash equilibrium and myopic strategies. Right panel shows the total pumping histories.

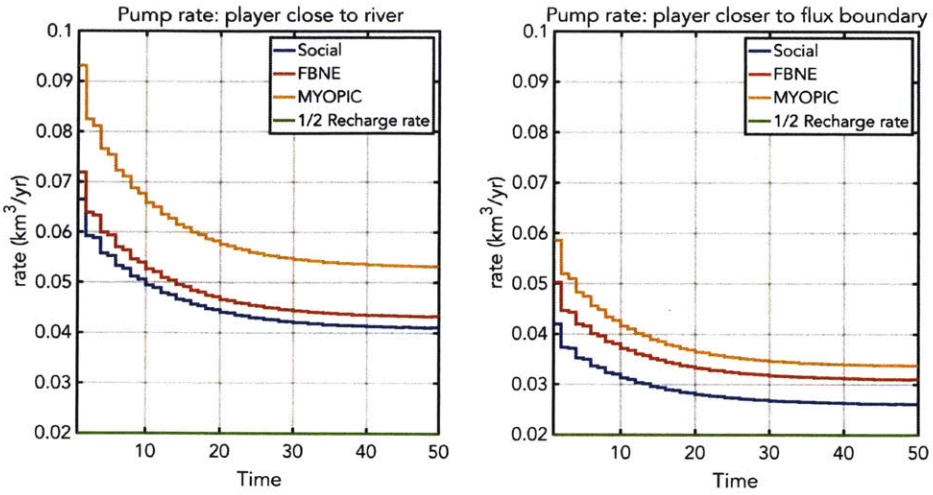

Figure 4-9: Individual player pump rates with an asymmetric well configuration under three decision strategies. Left panel shows pump rate of the well closer to the river boundary condition while the right panel shows the pump history of the well closer to the flux boundary.

optimum strategy. Table 4.2 summarizes the revenue and steady state pumping rates for the three strategies for the symmetric and asymmetric well configurations. The feedback Nash equilibrium and the myopic strategies are, as expected, more aggressive than the social optimum in drawing down the groundwater stock.

The spatially distributed model provides the capability to examine the effect of the geometric configuration in the player wells and reveals important asymmetries that are missed in a spatially lumped analysis of the groundwater common pool problem. Although the players in this example have the same benefit and cost functions, their total revenues differ for all decision strategies. The difference in their well locations creates about 16% change in the revenue generated. A lumped model fails to distinguish these two cases since it only considers the spatially averaged aquifer head in the pumping region.

4.3.3 Asymmetrical player benefits

The benefit function of the players in the nominal case generates $50000\$km^{-2}yr^{-1}$ using about $800mmyr^{-1}$ of water. We now consider a case where one of the players grows a crop that can generate 20% more revenue per unit area than the nominal case. The other player grows the nominal crop. The spatial configuration is the symmetrical arrangement considered earlier. Figure 4-10 shows the head contour curves and the flow lines at the steady state using a feedback Nash equilibrium strategy. The lower well is operated by the player with the higher value crop. This player pumps more and has a higher local drawdown than the player with the lower value crop. Figure 4-11 shows the variation of hydraulic head along the y-direction at $x = 10km$ while Figure 4-12 shows the pumping histories of the individual players. The steady-state

Case	**Revenue** (NPV) ($\times 10^6$\$)			**Steady state pump rate** (km^3yr^{-1})		
	Social	FBNE	Myopic	Social	FBNE	Myopic
Symmetric location	100.73	98.52 2.2%	91.93 8.7%	0.0732	0.0800 9.3%	0.0933 27.5%
Asymmetric location	84.6	82.7 2.5%	76.7 10%	0.067	0.074 10.4%	0.087 29.8%
Symmetric location Asymmetric benefits	132.16	129.01 2.4%	122.03 7.7%	0.0817	0.0888 8.7%	0.1029 26%

Table 4.2: Performance assessment of non-optimal management of the groundwater system w.r.t social optimum with a spatially distributed system. Percent values indicate relative difference to the social optimum.

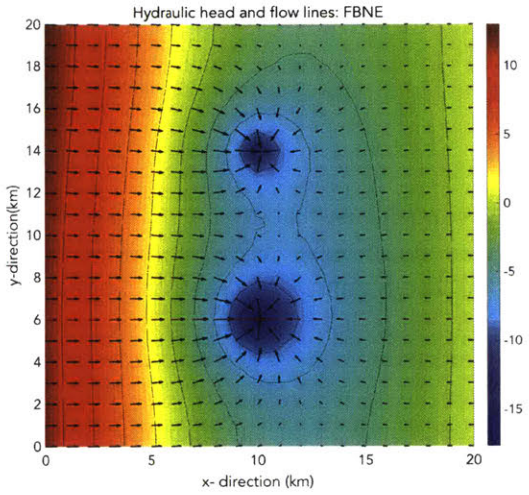

Figure 4-10: Contour plot and flow lines for a symmetric pump geometry but with asymmetry in the net benefit function using a feedback Nash equilibrium strategy

pumping rate for the player with the higher value crop is 20% higher. Table 4.2 quantifies the revenue inefficiency and pumpage increase of the non-cooperative and myopic strategies with respect to social optimum.

Overall, Table 4.2 indicates that the inefficiencies associated with a myopic strategy are greater than those obtained with an uncooperative Nash equilibrium that considers the future impacts of pumping. The inefficiency resulting from adoption of an uncooperative Nash strategy rather than a social optimum strategy is minor. Our spatially distributed analysis suggests that the local impacts of depletion at a player's well moderate pumping behavior sufficiently to diminish the adverse impacts of non-cooperative behavior, especially when future consequences are taken into account.

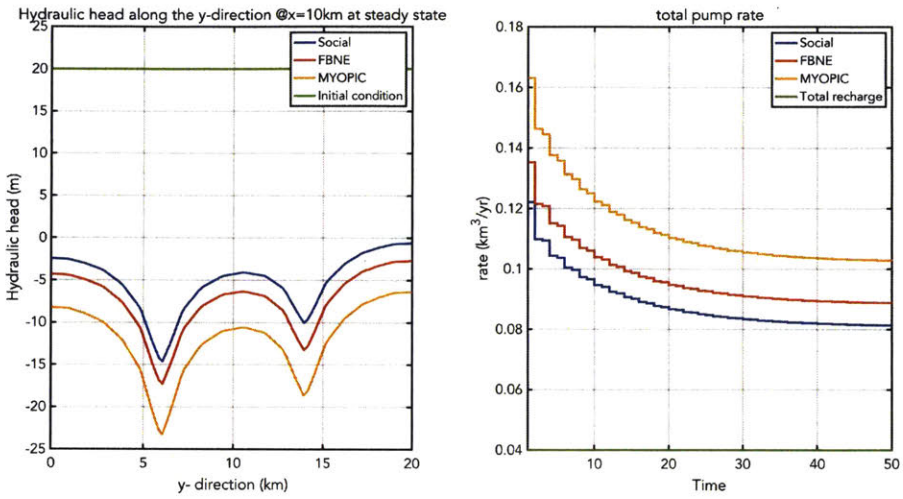

Figure 4-11: Left panel compares the aquifer head profile at steady state along the y direction at $x = 10km$ between social optimum, feedback Nash equilibrium and myopic strategies for asymmetric players. At $y = 6km$ the pump with the higher marginal benefit has greater drawdown than the other pumping location. Right panel shows the total pumping histories.

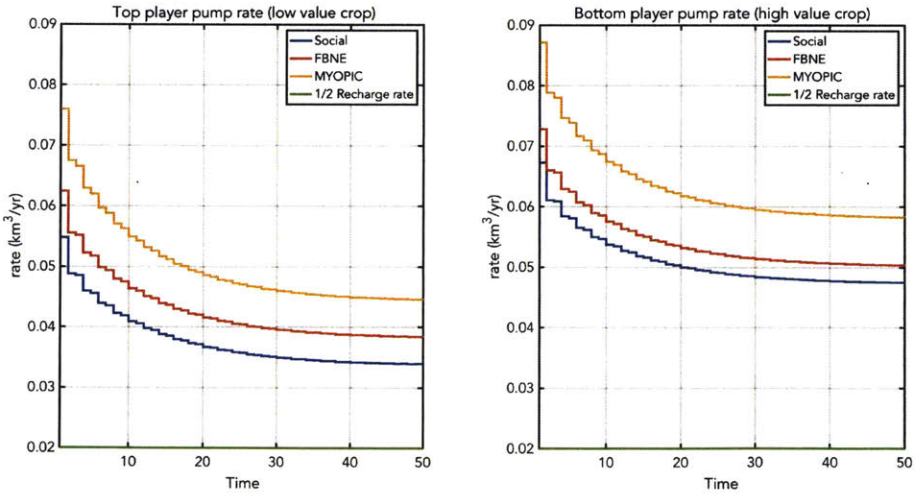

Figure 4-12: Individual player pump rates with an asymmetric well configuration under three decision strategies. Left panel shows pump rate of the well closer to the river boundary condition while the right panel shows the pump history of the well closer to the flux boundary.

4.4 Conclusions

This chapter extends our discussion of the common pool groundwater depletion problem by considering local spatial effects not included in the lumped analysis of Chapter 3. The spatially distributed description of groundwater flow used in this chapter provides a more realistic way to examine interactions between different players exploiting a common groundwater resource. With this description players decide on pumping rates based on local rather than spatially averaged well drawdowns. With suitable assumptions, spatially distributed groundwater dynamics can be reduced to set of affine state equations that adequately capture fine scale interactions between individual wells and between wells and sources of groundwater. When combined with a quadratic objective this affine formulation makes it possible to compute convenient closed form solutions for both uncooperative and cooperative aquifer management strategies and to readily evaluate the magnitude of economic inefficiencies resulting from excessive depletion.

The non-cooperative Nash equilibrium and myopic strategies both perform somewhat worse than the social optimum in terms of revenue generated, although the magnitude of the inefficiency is not large (about $2 - 10\%$). Our analysis shows that apart from the choice of the decision strategy, factors such as the the placement of pumping wells and their proximity to boundary sources play important roles in pumping behavior. Asymmetric pumping behavior between the players can arise in at least two ways – first, when players have different marginal benefit and cost functions and, second, when players have unequal access to water (even if they have symmetric objectives). Players with easier access to water tend to pump more than those who do not. A lumped model approach fails to capture the unequal access effect. The spatially distributed approach outlined here extends the scope of game-oriented research to capture key hydrologic features that have not been widely considered in the past. Consideration of these effects helps to identify how different problem specifications can influence the impact of uncooperative behavior on economic efficiency.

Chapter 5

Conclusions and scope for future work

This research presents novel approaches for better assessment in dynamic management of water resources. Using techniques from numerical programming and optimal control theory, two cases of water resources management are studied. The first is a hydropower reservoir operations problem and the second is a groundwater common pool resource problem. The underlying feature in these two cases is the dynamics of real-time decision making, either by the hydropower reservoir operator or the farmers sharing the aquifer. Making decisions dynamically gives one an ability to adapt to different external forcing such as uncertainty in the climate conditions or multiple agents competitively utilizing the water for their personal objective.

In the case of single agent hydropower operations problem, a coupled framework of fixing optimum contract and computing optimal real-time decision rule is developed. By using an ensemble based stochastic optimization approach, a decision-maker can obtain a probabilistic performance measure of their optimal contract across multiple objectives and across different predictive or non-predictive strategies before the operations begin. In addition, s/he can generate real-time revenue behavior and release decisions computed over different random future scenarios that carry the uncertainty information with different operational strategies. This conceptual framework is a useful tool to combine financial risk uncertainty to reservoir operational strategy. A possible direction to extend the scope of such coupled analysis is to incorporate random energy pricing in a multi-supplier and multi-buyer energy framework.

The work on the groundwater common pool resource problem improves upon the current understanding of the nature of groundwater sharing/ competition. With the timescales associated with groundwater flow and development in the agriculture sector, we observe that a transient analysis is necessary when modeling groundwater depletion. However, we find that there is marginal revenue difference ($\sim 5\%$) between non-cooperative strategies and the social optimum when applied across two different mathematical abstractions of the groundwater flow. The lumped model description, which uses a spatially averaged head to depict depletion is a lower-dimensional model than the spatially distributed approach. Although the two models give similar values of revenue efficiency, the difference lies in their depiction of flow behavior and water balance across the different source (river, recharge) and sink terms (pumping wells). The spatially distributed model shows local drawdowns at well locations and demotes

the exacerbated effect of strategic externality as observed in a lumped model. This is because, for a player to obtain groundwater access from distant regions, requires high local drawdown which increases his/ her pumping cost. Hence the estimate of strategic exernality or competition is dependent of various factors such a well configuration, local hydrology and objective criteria. The framework described in the study is flexible to capture different scenarios for a better assessment of groundwater management. A natural extension of this work, is to incorporate a player's choice to pump more wells and increase the cultivated area. This addition better signifies the development of mechanization in agriculture and large scale farms from subsidence level small scale farms. Also, the current analysis assumes a perfect knowledge of various parameters used in the system including the hydrology. However this is often not true. A stochastic treatment to the spatially distributed model may provide a better picture of expected utility from a shared aquifer system.

Appendix A

Hydropower optimization model: non dimensional form

Specifying the equivalent objective function, state equations ans inequality constraints with the non-dimensional variables

$$
\begin{aligned}
J'_{\mu'_{0:K-1}} &= \mathrm{E}\{\sum_{k=0}^{K-1}(1+r)^{-}k[g'[E'_k, E'_c] - \alpha'_Z Z'_k] + g'_K\} && \text{(A.1)}\\
g'(E'_k, E'_c) &= \alpha_1(E'_k - E'_c) + E'_c \;\; if \;\; E'_k \leq E'_c && \text{(A.2)}\\
g'(E'_k, E'_c) &= \alpha_2(E'_k - E'_c) + E'_c \;\; if \;\; E'_k > E'_c \\
\psi_k &= \log(I'_k) && \text{(A.3)}\\
\psi_{k+1} &= f_2(x_k, u_k, \omega_k) = \rho_\psi \psi_k + \omega_k && \text{(A.4)}\\
S'_{k+1} &= S'_k + \frac{1}{\tau_{res}}[I'_{k+1} - u'_k] - Z'_k && \text{(A.5)}\\
&= S_k + \frac{1}{\tau_{res}}[\exp \rho_\psi \psi_k + \omega_k - u'_k] - Z'_k \\
Z'_k &= \max\{S'_k + \frac{1}{\tau_{res}}[I'_k - u'_k] - 1, 0\} && \text{(A.6)}\\
E'_k &= \phi'(u'_k, h'_k, h'_{k+1}) = \frac{1}{E_{max}}\phi(u_{max}u'_k, h_{max}h'_k, h_{max}h'_{k+1}) && \text{(A.7)}\\
h'_k &= H'(S'_k) = \frac{1}{h_{max}}H(S_{max}S'_k) && \text{(A.8)}\\
u'_k &\leq u'_{max} \;\; for \; k = 0 : K-1 && \text{(A.9)}
\end{aligned}
$$

Appendix B

Groundwater Lumped model appendix

B.1 Definition of objective function and constraint matrices for the two-compartment dynamic game problem

Each agent i wants to maximize their objective L_i subject to the state equation implemented as an equality constraint

$$L_i = x_K^T Q_{xx,K} x_K + \sum_{k=1}^{K-1} (1+r)^{-k} \left[{u_k^i}^T Q_{uu,k} u_k^i + 2 x_k^T Q_{ux,k}^i u_k^i + 2 S_{u,k}^i \right] \quad \text{(B.1)}$$

$$\forall i \in N, \ @k = 1, x_1 \ is \ given \quad \text{(B.2)}$$

$$x_{k+1} = F_k x_k + \sum_{i=1}^{N} G_k^i u_k^i + E_k \quad \text{(B.3)}$$

where $x_k = [h_{Ok}, h_{Ik}]^T$ is the aquifer head in the outer and the inner aquifer;

$$F_k = \begin{bmatrix} 1 - \frac{\lambda_{OI}\Delta t}{A_O \theta} & \frac{\lambda_{OI}\Delta t}{A_O \theta} \\ \frac{\lambda_{IR}\Delta t}{A_I \theta} & 1 - \frac{(\lambda_{OI}+\lambda_{IR})\Delta t}{A_I \theta} \end{bmatrix} \quad \text{(B.4)}$$

$$G_k^i = \begin{bmatrix} 0 \\ -\frac{\Delta t}{A_i \theta} \end{bmatrix} \quad \text{(B.5)}$$

$$E_k = \begin{bmatrix} \frac{R}{A_O \theta} \\ \frac{\Delta t \lambda_{IR} h_R}{A_i \theta} \end{bmatrix} \quad \text{(B.6)}$$

$$Q_{uu,k}^i = b_2^i; \ \ Q_{ux,k}^i = [0 \ 0.5c\eta]; \ S_{u,k}^i = b_1^i - 0.5c\eta h_g, \ \forall i \in N \quad \text{(B.7)}$$

The present analysis assumes the parameter based matrices, $[F, G^i, E, Q_{uu}^i, Q_{ux}^i, S_{ux}^i]$, $\forall i \in N$ to be constant in time but can easily be generalized to a time varying behavior

too.

B.2 Social optimum decision strategy

In this strategy a single player decides the pumping rate for the all the wells. Therefore the single agent has a 2-state N-control dynamic problem to solve. The N objectives from the agents are augmented together to give a modified linear quadratic objective L_i and a modified affine state equation as an equality constraint.

$$L = x_K^T Q_{xx,K} x_K + \sum_{k=1}^{K-1} (1+r)^{-k} \left[u_k^T Q_{uu} u_k + 2x_k^T Q_{ux} u_k + 2S_u u_k \right] \quad \text{(B.8)}$$

$$\textit{subject to} \quad \text{(B.9)}$$

$$x_{k+1} = Fx_k + Gu_k + E_k;\ @k = 1;\ x_1 \textit{ is given} \quad \text{(B.10)}$$

where

$$\begin{aligned} u_k &= [u_k^1, u_k^2, , u_k^i, ., u_k^N]^T \\ Q_{xx,K} &= \sum_{i=1}^{N} Q_{xx,K}^i \\ Q_{uu} &: \textit{ diagonal } N \times N \textit{ matrix with } Q_{uu}^i \textit{ as diagonal elements} \\ Q_{ux} &= \left[Q_{ux}^1, Q_{ux}^2, , Q_{ux}^i, , Q_{ux}^N \right];\ 2 \times N \textit{ matrix} \\ S_u &= [S_u^1, S_u^2, S_u^i, , S_u^N];\ 1 \times N \textit{ matrix} \\ G &= [G^1, G^2, G^i, , G^N];\ 2 \times N \textit{ matrix} \end{aligned}$$

The above equations transform to a linear quadratic control problem with affine constraint. This can be solved using dynamic programming that recursively solves from the final time step and generates a decision strategy for each time step. Due the special structure of this problem, the decision strategy also holds an affine structure. From the final time step,

$$V_{K-1} = \max_{u_{K-1}} u_{K-1}^T Q_{uu} u_{K-1} + 2x_{K-1}^T Q_{ux} u_{K-1} + 2S_u u_{K-1} + x_K^T \bar{Q}_{xx,K} x_K \quad \text{(B.11)}$$

$$\bar{Q}_{xx,K} = (1+r)^{-1} Q_{xx,K} \quad \text{(B.12)}$$

Substituting, $x_K = Fx_{K-1} + Gu_{K-1} + E$ and applying the first order condition w.r.t to u_{K-1}.

$$\begin{aligned} \frac{\partial V_{K-1}}{\partial u_{K-1}} &= 0 = Q_{uu} u_{K-1} + Q_{ux}^T x_{K-1} + S_u + G^T \bar{Q}_{xx,K} (Fx_K + Gu_K + E) \\ u_{SO,K-1} &= -\left(Q_{uu} + G^T \bar{Q}_{xx,K} G \right)^{-1} \left((Q_{ux}^T + G^T \bar{Q}_{xx,K} F) x_{K-1} + G^T \bar{Q}_{xx,K} E \right) \\ u_{SO,K-1}^i &= \gamma_{S),K-1}^i (x_{K-1}) = \Gamma_{SO,K}^i x_{K-1} + \Upsilon_{SO,K-1}^i \end{aligned}$$

As mentioned the u_{K-1} takes an affine structure with respect to x_{K-1}. This process is repeated backwards in time from $k = K-1$ to $k = 1$. The value function V_{K-1} always obtained by substituting $\gamma^i_{SO,K-1}(x_{K-1})$

$$V_{K-1}(x_{K-1}) = x^T_{K-1} P_{K-1} x_{K-1} + 2L_{K-1} x_{K-1} + \epsilon_{K-1}$$

Therefore at any stage k,

$$x_k^T P_k x_k + 2L_k x_k = \max_{u_k} u_k^T Q_{uu} u_k + 2x_k^T Q_{ux} u_k + 2S_u u_k + (1+r)^{-1}(x_{k+1}^T P_{k+1} x_{k+1} + 2L_{k+1} x_{k+1}) \tag{B.13}$$

$$\begin{aligned} \gamma_{SO,k} &= -(Q_u u + G^T P_{xx,k+1} G)^{-1} \left((Q_u x^T + G^T P_{xx,k+1} F) x_k + G^T P_{xx,k+1} E + \bar{L}_{k+1} G \right) \\ P_K &= Q_{xx,K},\ L_K = [0,,0]^T \end{aligned}$$

P_k and L_k are recursively generated by substituting $\gamma_(SO,k)$ in B.13.

B.3 Open loop Nash equilibrium strategy

In an open loop strategy, the players commit to (uncooperatively) to best personal interest from the initial stage ($k = 1$) based on the perfect information of the initial state. Therefore, we need to compute a Nash equilibrium strategy with the entire present value of the net benefit function for each agent. To compute this equilibrium, the dynamic problem is converted to an equivalent static problem for each agent. The objective function of each agent i can be written as

$$L_i = X^T \underline{Q}^i_{xx} X + \underline{u}^{iT} \underline{Q}^i_{uu} \underline{u}^i + 2X^T \underline{Q}^i_{ux} \underline{u}^i \tag{B.14}$$

subject to

$$X = \underline{F} X_1 + \sum_{i=1}^{N} \underline{G}^i \underline{u}^i + \underline{E} \tag{B.15}$$

where $X, X_1, \underline{u}^i$ are augmented vectors

$$X = \left[x_1^T | ... | x_K^T\right]^T \tag{B.16}$$

$$X_1 = \left[x_1^T | ... | x_1^T\right]^T \tag{B.17}$$

Similarly the matrices $\underline{F}.\underline{G}^i, \underline{E}, \underline{Q}^i_{uu}, \underline{Q}^i_{ux}, \underline{S}^i_u$ are appropriately generated using the state equation 3.4 ([7]). With these definitions the open loop problem can be solved as a static game problem.

Substituting equation B.15 in B.14 and applying the first order condition om L_i

respect to $\underline{u}^i$ gives a set of N coupled linear matrix equations in $\underline{u}^i$.

$$\begin{aligned}\frac{\partial L_i}{\partial \underline{u}^i} = 0 \quad &= \underline{G}^{iT}\underline{Q}^i_{xx}\left(\underline{F}X_1 + \underline{G}^i\underline{u}^i + \sum_{j\neq i}\underline{G}^j\underline{u}^i + E\right) + 2\underline{Q}^i_{uu}\underline{u}^i \\ &+2\underline{G}^{iT}\underline{Q}^i_{ux}\underline{G}^i\underline{u}^i + \underline{Q}^{i\ T}_{ux}\left(\underline{F}X_1 + \underline{G}^i\underline{u}^i + \sum_{j\neq i}\underline{G}^j\underline{u}^i + E\right) \qquad \text{(B.18)} \\ &+\underline{S}^i_u \quad \forall i \in N \\ \underline{u}^i \quad = \quad &P^i(\underline{G}^{iT}\underline{Q}^i_{xx}\underline{F} + \underline{Q}^i_{ux})\left(X_1 + \sum_{j\neq i}\underline{G}^j\underline{u}^j)\right) \\ &+P^i\left(\underline{G}^{iT}\underline{Q}^i_{xx}\underline{E} + \underline{Q}^i_{ux}\underline{E} + \underline{S}^i_u\right) \qquad \text{(B.19)}\end{aligned}$$

where $P^i = -\left(\underline{G}^{iT}\underline{Q}^i_{xx}\underline{G}^i + \underline{Q}^i_{uu} + \underline{G}^i\underline{Q}^i_{ux}\underline{G}^i\right)^{-1}$

Solving this system of equations gives an affine solution for $\underline{u}^i$

$$\left[u^i_{OL,1}, ..., u^i_{OL,K-1}\right] = \underline{u}^{i*} = \gamma^i_{OL}(X_1) = \Gamma^i_{OL}X_1 + \upsilon^i_{OL} \qquad \text{(B.20)}$$

The solution is $K-1$ length vector that is derived only from the initial state x_1

B.4 Closed loop/ Feedback Nash equilibrium decision strategy (FBNE)

In this strategy, at each stage the players maximize their future benefits for any possible state of the system to satisfy the condition of subgame perfect Nash equilibrium. Unlike the open-loop solution that generates an action strategy at the initial stage, the feedback NE generates a decision rule for each player at each stage of the game. It is obtained by dynamic programming that recursively solves a set of N-coupled matrix linear equations. Starting from the stage $K-1, \forall i \in N$

$$V_{i,K-1} = \max_{u^i_{K-1}} u^i_{K-1}Q^i_{uu}u^i_{K-1} + 2S^i_u u^i_{K-1} + x^T_K Q_{xx}x_K, \ \forall i \in N \qquad \text{(B.21)}$$

subject to equation B.3 (State equation)

Similar to the social optimum, the value function $V_{i,K-1}$ is maximized with respect to the control $u_{i,K-1}$ (by applying the first order conditions). This generated a system of linear set of equations with respect to $u_{i,K-1}$

$$\frac{\partial V_{i,K-1}}{\partial u^i_{K-1}} = 0 = Q^i_{uu} u^i_{K-1} + {Q^i_{ux}}^T x_{K-1} + S^i_u$$

$$G^{iT} \bar{Q}^i_{xx,K} \left(F x_K + G^i u^i_{K-1} + \sum_{j \neq i} G^j u^j_{K-1} + E \right)$$

$$u^i_{K-1} = Y^i_{FB,K-1}((Q^i_{ux}{}^T + G^{iT} \bar{Q}^i_{xx} + F) x_K + G^{iT} \bar{Q}^i_{xx,K} \sum_{j \neq i} G^j u^j_{K-1} \quad \text{(B.22)}$$

$$+ G^{iT} \bar{Q}^i_{xx,K} E + S^i_u) \quad \text{(B.23)}$$

where $Y^i_{FB,K-1} = -(Q^i_{uu} + G^{iT} \bar{Q}^i_{xx,K} G^i)^{-1}$

Solving the N set coupled linear equations with respect to u^i_{K-1} gives an affine decision rule.

$$u^i_{FB,K-1} = \gamma^i_{FB,K-1}(x_{K-1}) = \Gamma^i_{FB,K-1} x_{K-1} + \Upsilon^i_{FB,K-1} \quad \text{(B.24)}$$

Substituting the optimal strategy at stage $K-1$ in equation B.4 yields a quadratic value function V^i_{K-1}, $\forall i \in N$

$$V^i_{K-1} = x^T_{K-1} P^i_{K-1} x_{K-1} + 2 L^i_{K-1} x_{K-1} + \epsilon^i_{K-1} \quad \text{(B.25)}$$

Repeating this recursively generates a subgame perfect equilibrium solution at each time step. At any stage k,

$$\begin{aligned} x^T_k P^i_k x_k + 2 L^i_k x_k = & \max_{u^i_k} {u^i_k}^T Q^i_{uu} u^i_k + 2 x^T_k Q^i_{ux} u^i_k + 2 S^i_u u_k \\ & (1+r)^{-1} (x^T_{k+1} P^i_{k+1} x_{k+1} + 2 L^i_{k+1} x_{k+1}) \end{aligned} \quad \text{(B.26)}$$

$$u^i_{FB,k} = \gamma^i_{FB,k}(x_k) = \Gamma^i_{FB,k} x_k + \Upsilon^i_{FB,k} \quad \text{(B.27)}$$

B.5 Myopic Strategy

Under a myopic strategy a player maximizes current net benefits only, without regard for future benefits or the strategic externality of other players. This is easily computed each stage by equating the current marginal revenue to the marginal cost. For each player $i \in N$

$$V_{i,k}(x_k) = \max_{u^i_k} {u^i_k}^T Q^i_{uu} u^i_k + 2 x^T_k Q^i_{ux} u^i_k + 2 S^i_u u^i_k \quad \text{(B.28)}$$

$$\frac{\partial V_{i,k}}{\partial u_{i,k}} = 0 = Q^i_{uu} u^i_k + Q^i_{ux} x_k + S^i_u \quad \text{(B.29)}$$

$$u^i_k = -{Q^i_{uu}}^{-1} ({Q^i_{ux}}^T x_k + S^i_u) \quad \text{(B.30)}$$

$$u^i_{MY,k} = \gamma^i_{MY,k}(x_k) = \Gamma^i_{MY,k} x_k + \Upsilon^i_{MY,k} \quad \text{(B.31)}$$

Appendix C

Spatially distributed groundwater system

The coefficient matrix in equation 4.6 are obtained from finite element based discretization scheme. For a standard elliptic equation:

$$d\frac{\partial h}{\partial t} - \nabla \cdot (c \cdot \nabla h) = f \tag{C.1}$$

This equation holds in constrained domain Ω. The boundary conditions can be of the form Neumann boundary (flux-boundary) or a Dirchlet boundary (Head-boundary) condition.

$$Dirchlet: \quad \alpha h = r, \ on \ the \ boundary \ \partial\Omega \tag{C.2}$$

$$Neumann: \quad \overrightarrow{n} \cdot (\nabla \cdot h) + qh = g \ on \ \partial\Omega \tag{C.3}$$

$\overrightarrow{n}$ is the outward unit normal. g, q, α, r are functions defined on $\partial\Omega$.

Assume that h is a solution of the PDE system. Multiply the entire equation by v is an arbitrary test function and intergrte over the domain Ω. For a steady state analysis $d = 0$,

$$\int_\Omega -(\nabla \cdot c \cdot h)v \ dx = \int_\Omega fv \ dx \tag{C.4}$$

Integrate by parts (using Green's formula) to obtain:

$$\int_\Omega ((c\nabla h) \cdot \nabla v)dx - \int_{\partial\Omega} \overrightarrow{n} \cdot (c\nabla \cdot h)v \ ds = \int_\Omega fv \ dx \tag{C.5}$$

Adding the boundary condition to replace the integral at the boundary

$$\int_\Omega ((c\nabla h) \cdot \nabla v)dx - \int_{\partial\Omega} (-qh + g)v \ ds = \int_\Omega fv \ dx \tag{C.6}$$

By projecting the solution u and v that belong to an function space V, the above equation projected unto a finite dimensional subspace $V_{N_p} \subset V$. As the differential

operator is linear, the variational equation must solve must satisy N_p test functions $\phi_i \in V_{N_p}$.

$$\int_\Omega ((c\nabla h) \cdot \nabla phi_i - f\phi_i)dx - \int_{\partial\Omega} (-qh + g)\phi_i \; ds = 0, \; for \; i = 1, 2.., N_p \tag{C.7}$$

Expressing the solution h in the same basis of V_{N_p}

$$h(x) = \sum_{j=1}^{N_p} H_j \phi_j(x) \tag{C.8}$$

$$\sum_{j=1}^{N_p} ((c\nabla\phi_j) \cdot \nabla\phi_i dx + \int_{\partial\Omega} q\phi_j \; ds)U_j = \int_\Omega f\phi_i dx + \int_{\partial\Omega} g\phi_i ds, \; for \; i = 1, .., N_p \tag{C.9}$$

$$(K + M + Q)H = F + G \tag{C.10}$$

where

K : stiffness matrix: $(c\nabla\phi_j) \cdot \nabla\phi_i dx$
M : mass matrix : integral of d coefficient against basis function
Q: integral of q coefficient against basis function
F: integral of f coefficient against basis function

These coefficient matrices are then used to construct the corresponding low-dimensional spatially distributed system dynamical system.

$$M(H_{k+1} - H_k) + 0.5K(H_{k+1} + H_k) = F_{river}H_r + F_{flux}R + \sum_{i=1}^{N} F_{pump_i}Q_i \tag{C.11}$$

$$(M + 0.5K)\, H_{k+1} = (M - 0.5K)\, H_k + F_{river}H_r + F_{flux}R + \sum_{i=1}^{N} F_{pump_i}Q_i \tag{C.12}$$

Bibliography

[1] H Abou-Kandil, G Freiling, V Ionescu, and G Jank. *Matrix Riccati Equations in Control and Systems Theory*, volume 49. 2003.

[2] Werner Aeschbach-Hertig and Tom Gleeson. Regional strategies for the accelerating global problem of groundwater depletion. *Nature Geoscience*, 5(12):853–861, 2012.

[3] Richard C. Allen and Micha Gisser. Competition Versus Optimal Control in Groundwater Pumping When Demand is Nonlinear. *Water Resources Research*, 20(7):752–756, aug 1984.

[4] Yaakov Bar-Shalom and Edison Tse. Dual Effect, Certainty Equivalence, and Separation in Stochastic Control. *IEEE Transactions on Automatic Control*, 19(5):494–500, 1974.

[5] Toni Barjas Blanco, Patrick Willems, Po Kuan Chiang, Niels Haverbeke, Jean Berlamont, and Bart De Moor. Flood regulation using nonlinear model predictive control. *Control Engineering Practice*, 18(10):1147–1157, oct 2010.

[6] Mario T. L. Barros, Frank T-C. Tsai, Shu-li Yang, Joao E.G. Lopes, and William W-G. Yeh. Optimization of Large-Scale Hydropower System Operations. *Journal of Water Resources Planning and Management*, 129(June):11, 2003.

[7] T. Basar. A counterexample in linear-quadratic games: Existence of nonlinear nash solutions. *Journal of Optimization Theory and Applications*, 14(4):425–430, oct 1974.

[8] Tamer Basar and Geer Jan Olsder. *Dynamic Noncooperative Game Theory*. 1999.

[9] J Bear. *Hydraulics of groundwater*. 2012.

[10] R. Bellman. Dynamic Programming and Lagrange Multipliers. *Proceedings of the National Academy of Sciences*, 42(10):767–769, oct 1956.

[11] Dimitri P. Bertsekas and David A. Castañon. Rollout algorithms for stochastic scheduling problems. *Journal of Heuristics*, 5(1):89–108, 1999.

[12] DP Bertsekas. *Dynamic programming and optimal control*, volume II. 1995.

[13] John D. Bredehoeft. The water budget myth revisited: Why hydrogeologists model, 2002.

[14] Nicholas Brozović, David L. Sunding, and David Zilberman. On the spatial nature of the groundwater pumping externality. *Resource and Energy Economics*, 32(2):154–164, 2010.

[15] A. Castelletti, D. de Rigo, A. E. Rizzoli, R. Soncini-Sessa, and E. Weber. Neuro-dynamic programming for designing water reservoir network management policies. *Control Engineering Practice*, 15(8):1031–1038, aug 2007.

[16] Jps P S Catalão, H M I Pousinho, and J Contreras. Hydro generation scheduling and offering strategies considering price uncertainty and risk management. *Proceedings of the 17th PSCC, Stockholm, Sweden*, pages 1–24, 2011.

[17] Cristiano Cervellera and Marco Muselli. Efficient sampling in approximate dynamic programming algorithms. *Computational Optimization and Applications*, 38(3):417–443, jun 2007.

[18] J. C. Engwerda and Salmah. Necessary and Sufficient Conditions for Feedback Nash Equilibria for the Affine-Quadratic Differential Game. *Journal of Optimization Theory and Applications*, 157(2):552–563, 2013.

[19] Jacob Engwerda. Feedback Nash equilibria in the scalar infinite horizon LQ-game. *Automatica*, 36(1):135–139, 2000.

[20] J.C. Engwerda. Computational aspects of the open-loop Nash equilibrium in linear quadratic games. *Journal of Economic Dynamics and Control*, 22(8-9):1487–1506, 1998.

[21] J. S. Famiglietti. The global groundwater crisis, 2014.

[22] A. K. Fink. Tehri hydro power complex on the Bhagirathi river in India. *Hydrotechnical Construction*, 34(8-9):479–484, 2000.

[23] Benjamin T. Foster, Jordan D. Kern, and Gregory W. Characklis. Mitigating hydrologic financial risk in hydropower generation using index-based financial instruments. *Water Resources and Economics*, 10:45–67, 2015.

[24] G Freiling, G Jank, and H Aboukandil. On Global Existence of Solutions to Coupled Matrix Riccati-Equations in Closed-Loop Nash Games. *IEEE Trans. Automat. Contr.*, 41(2):264–269, 1996.

[25] Devin L. Galloway and Thomas J. Burbey. Review: Regional land subsidence accompanying groundwater extraction. *Hydrogeology Journal*, 19(8):1459–1486, 2011.

[26] Carlos E. García, David M. Prett, and Manfred Morari. Model predictive control: Theory and practice-A survey. *Automatica*, 25(3):335–348, 1989.

[27] Míriam R. García, Carlos Vilas, Lino O. Santos, and Antonio A. Alonso. A robust multi-model predictive controller for distributed parameter systems. *Journal of Process Control*, 22(1):60–71, 2012.

[28] Roy Gardner, Michael R. Moore, and James M. Walker. Governing a groundwater commons: A strategic and laboratory analysis of western water law. *Economic Inquiry*, 35(2):218–234, apr 1997.

[29] M. Giuliani and A. Castelletti. Assessing the value of cooperation and information exchange in large water resources systems by agent-based optimization. *Water Resources Research*, 49(7):3912–3926, 2013.

[30] Steven M. Gorelick and Irwin Remson. Optimal dynamic management of groundwater pollutant sources. *Water Resources Research*, 18(1):71–76, feb 1982.

[31] Steven M. Gorelick and Chunmiao Zheng. Global change and the groundwater management challenge. *Water Resources Research*, 51(5):3031–3051, may 2015.

[32] Warren A. Hall, William S. Butcher, and Austin Esogbue. Optimization of the Operation of a MultipleâĂŘPurpose Reservoir by Dynamic Programming. *Water Resources Research*, 4(3):471–477, 1968.

[33] Eric R. Hooper, Aris P. Georgakakos, and Dennis P. Lettenniaier. Optimal Stochastic Operation of Salt River Project, Arizona. *Journal of Water Resources Planning and Management*, 117(5):566–587, 1991.

[34] F O R Irrigation. San Joaquin Valley Grower Irrigation. Technical Report December, 2002.

[35] Sharon A. Johnson, Jery R. Stedinger, Christine A. Shoemaker, Ying Li, and José Alberto Tejada-Guibert. Numerical Solution of Continuous-State Dynamic Programs Using Linear and Spline Interpolation. *Operations Research*, 41(3):484–500, 1993.

[36] C. S. Kim, Michael R. Moore, John J. Hanchar, and Michael Nieswiadomy. A dynamic model of adaptation to resource depletion: theory and an application to groundwater mining. *Journal of Environmental Economics and Management*, 17(1):66–82, jul 1989.

[37] Keith C. Knapp and Lars J. Olson. The economics of conjunctive groundwater management with stochastic surface supplies. *Journal of Environmental Economics and Management*, 28(3):340–356, 1995.

[38] Leonard F. Konikow and Eloise Kendy. Groundwater depletion: A global problem. *Hydrogeology Journal*, 13(1):317–320, 2005.

[39] Phoebe Koundouri. Current issues in the economics of groundwater resource management. *Journal of Economic Surveys*, 18(5):703–740, 2004.

[40] Phoebe Koundouri. Potential for groundwater management: Gisser-Sanchez effect reconsidered. *Water Resources Research*, 40(6):1–13, 2004.

[41] John W. Labadie. Optimal Operation of Multireservoir Systems: State-of-the-Art Review. *Journal of Water Resources Planning and Management*, 130(2):93–111, 2004.

[42] Jay H. Lee. Model predictive control: Review of the three decades of development. *International Journal of Control, Automation and Systems*, 9(3):415–424, jun 2011.

[43] Hartmut Linke. A model-predictive controller for optimal hydro-power utilization of river reservoirs. In *Proceedings of the IEEE International Conference on Control Applications*, pages 1868–1873. IEEE, sep 2010.

[44] Lennart Ljung. Estimating Linear Time-invariant Models of Nonlinear Time-varying Systems. *European Journal of Control*, 7(2-3):203–219, 2001.

[45] Hugo A. Loáiciga. Analytic game - Theoretic approach to ground-water extraction. *Journal of Hydrology*, 297(1-4):22–33, 2004.

[46] Kaveh Madani and Ariel Dinar. Exogenous regulatory institutions for sustainable common pool resource management: Application to groundwater. *Water Resources and Economics*, 2-3:57–76, 2013.

[47] D. Q. Mayne, J. B. Rawlings, C. V. Rao, and P. O.M. Scokaert. Constrained model predictive control: Stability and optimality. *Automatica*, 36(6):789–814, 2000.

[48] Birger Mo, Anders Gjelsvik, and Asbjørn Grundt. Integrated risk management of hydro power scheduling and contract management. *IEEE Transactions on Power Systems*, 16(2):216–221, may 2001.

[49] Marc F. Müller, Michèle C. Müller-Itten, and Steven M. Gorelick. How Jordan and Saudi Arabia are avoiding a tragedy of the commons over shared groundwater. *Water Resources Research*, 53(7):5451–5468, jul 2017.

[50] Kevin B. Mulligan, Casey Brown, Yi Chen E Yang, and David P. Ahlfeld. Assessing groundwater policy with coupled economic-groundwater hydrologic modeling. *Water Resources Research*, 50(3):2257–2274, 2014.

[51] Megumi Nakao, Dennis Wichelns, and John Montgomery. Game Theory Analysis of Competition for Groundwater Involving El Paso , Texas and Ciudad Juarez , Mexico Game Theory Analysis of Competition for Groundwater Involving El Paso , Texas and Ciudad Juarez , Mexico. In *Water*, pages 18–31, 2002.

[52] T. R. Neelakantan and N. V. Pundarikanthan. Hedging rule optimisation for water supply reservoirs system. *Water Resources Management*, 13(6):409–426, 1999.

[53] R R Negenborn, P J Van Overloop, T Keviczky, and B De Schutter. Distributed model predictive control of irrigation canals. *Network and Heterogeneous Media*, 4(2):359–380, 2009.

[54] Donald H. Negri. The common property aquifer as a differential game. *Water Resources Research*, 25(1):9–15, 1989.

[55] Edgar Perea-López, B. Erik Ydstie, and Ignacio E. Grossmann. A model predictive control strategy for supply chain optimization. In *Computers and Chemical Engineering*, volume 27, pages 1201–1218, 2003.

[56] Jeffrey M Peterson and Alexander E Saak. Spatial externalities in aquifers with varying thickness: Theory and numerical results for the Ogallala aquifer *. 2013.

[57] Lisa Pfeiffer and C. Y.Cynthia Lin. Groundwater pumping and spatial externalities in agriculture. *Journal of Environmental Economics and Management*, 64(1):16–30, jul 2012.

[58] Steven P Phillips, Karen R Burow, Diane L Rewis, Jennifer Shelton, and Bryant Jurgens. Hydrogeologic Setting and Ground-Water Flow Simulations of the San Joaquin Valley Regional Study Area, California Hydrogeologic Settings and Ground-Water Flow Simulations for Regional Studies of the Transport of Anthropogenic and Natural Contaminants to P.

[59] Bill Provencher and Oscar Burt. The externalities associated with the common property exploitation of groundwater. *Journal of Environmental Economics and Management*, 24(2):139–158, 1993.

[60] S.J Qin and T.A Badgwell. A survey of industrial model predictive control technology. *Control Engineering Practice 11*, pages 733 – 764, 2003.

[61] Deepti Rani and Maria Madalena Moreira. Simulation-optimization modeling: A survey and potential application in reservoir systems operation. *Water Resources Management*, 24(6):1107–1138, apr 2010.

[62] James B. Rawlings. Tutorial Overview of Model Predictive Control. *IEEE Control Systems*, 20(3):38–52, jun 2000.

[63] Santiago J. Rubio and Begoña Casino. Competitive versus efficient extraction of a common property resource: The groundwater case. *Journal of Economic Dynamics and Control*, 25(8):1117–1137, aug 2001.

[64] Santiago J. Rubio and Begoña Casino. Strategic Behavior and Efficiency in the Common Property Extraction of Groundwater. *Environmental and Resource Economics*, 26(1):73–87, 2003.

[65] Alexander E. Saak and Jeffrey M. Peterson. Groundwater use under incomplete information. *Journal of Environmental Economics and Management*, 54(2):214–228, 2007.

[66] Yahya Saleh, Ülkü Gürler, and Emre Berk. Centralized and decentralized management of groundwater with multiple users. *European Journal of Operational Research*, 215(1):244–256, 2011.

[67] Bridget R. Scanlon, Ian Jolly, Marios Sophocleous, and Lu Zhang. Global impacts of conversions from natural to agricultural ecosystems on water resources: Quantity versus quality. *Water Resources Research*, 43(3), mar 2007.

[68] Pepijn Schreinemachers and Thomas Berger. An agent-based simulation model of human-environment interactions in agricultural systems. *Environmental Modelling and Software*, 26(7):845–859, 2011.

[69] Stefan Siebert and Petra Döll. Quantifying blue and green virtual water contents in global crop production as well as potential production losses without irrigation. *Journal of Hydrology*, 384(3-4):198–217, apr 2010.

[70] Tobias Siegfried and Wolfgang Kinzelbach. A multiobjective discrete stochastic optimization approach to shared aquifer management: Methodology and application. *Water Resources Research*, 42(2), 2006.

[71] A. W. Starr and Y. C. Ho. Nonzero-sum differential games. *Journal of Optimization Theory and Applications*, 3(3):184–206, mar 1969.

[72] Jery R. Stedinger, Beth A. Faber, and Jonathan R. Lamontagne. Developments in Stochastic Dynamic Programming for Reservoir Operation Optimization. *World Environmental and Water Resources Congress 2013*, pages 1266–1278, 2013.

[73] Claudia M. Stickler, Michael T. Coe, Marcos H. Costa, Daniel C. Nepstad, David G. McGrath, Livia C. P. Dias, Hermann O. Rodrigues, and Britaldo S. Soares-Filho. Dependence of hydropower energy generation on forests in the Amazon Basin at local and regional scales. *Proceedings of the National Academy of Sciences*, 110(23):9601–9606, jun 2013.

[74] Jordan F. Suter, Joshua M. Duke, Kent D. Messer, and Holly A. Michael. Behavior in a spatially explicit groundwater resource: Evidence from the lab. *American Journal of Agricultural Economics*, 94(5):1094–1112, oct 2012.

[75] Ming-Yen Tu, Nien-Sheng Hsu, and William W.-G. Yeh. Optimization of Reservoir Management and Operation with Hedging Rules. *Journal of Water Resources Planning and Management*, 129(2):86–97, mar 2003.

[76] Peter Jules van Overloop, Steven Weijs, and Sjoerd Dijkstra. Multiple Model Predictive Control on a drainage canal system. *Control Engineering Practice*, 16(5):531–540, may 2008.

[77] Yoshihide Wada, Ludovicus P.H. Van Beek, Cheryl M. Van Kempen, Josef W.T.M. Reckman, Slavek Vasak, and Marc F.P. Bierkens. Global depletion of groundwater resources. *Geophysical Research Letters*, 37(20):n/a–n/a, oct 2010.

[78] Ralph A. Wurbs. Reservoir-System SImulations and Optimization Models. *Journal of Water Resources Planning and Management*, 119(4):455–472, 1993.

[79] Ralph A. Wurbs. ReservoirâĂŘSystem Simulation and Optimization Models. *Journal of Water Resources Planning and Management*, 119(4):455–472, jul 1993.

[80] S Yakowitz. Dynamic programming applications in water resources. *Water Resources Research*, 18(4):673–696, 1983.

[81] Yi Chen E Yang, Ximing Cai, and Dušan M. Stipanović. A decentralized optimization algorithm for multiagent system-based watershed management. *Water Resources Research*, 45(8), 2009.

[82] William WâĂŘG Yeh. Reservoir Management and Operations Models: A StateâĂŘofâĂŘtheâĂŘArt Review. *Water Resources Research*, 21(12):1797–1818, 1985.

[83] Jiing Yun You and Ximing Cai. Hedging rule for reservoir operations: 2. A numerical model. *Water Resources Research*, 44(1):n/a–n/a, jan 2008.

[84] Tongtiegang Zhao, Jianshi Zhao, and Dawen Yang. Improved Dynamic Programming for Hydropower Reservoir Operation. *Journal of Water Resources Planning and Management*, 140(3):365–374, 2014.