

Analytical Tools for Integrating Transfers into Water Resource Management Strategies

Brian R. Kirsch

A dissertation submitted to the faculty of the University of North Carolina at Chapel Hill
in partial fulfillment of the requirements for the degree of Doctor of Philosophy in the
Department of Environmental Science and Engineering

Chapel Hill

2010

Approved by:

Gregory W. Characklis, Ph.D.

Francis DiGiano, Ph.D.

Tim Kelley, Ph.D.

David Moreau, Ph.D.

Philip Singer, Ph.D.

ABSTRACT

Brian R. Kirsch

Analytical Tools for Integrating Transfers into Water Resource Management Strategies

(Under the direction of Gregory W. Characklis, Ph.D.)

Many municipalities within the United States anticipate rising demand for water as populations grow. Traditionally, rising demand has often been addressed via infrastructure projects, such as reservoirs. However, a variety of factors has combined to make such projects less attractive, such as increased development costs, stricter environmental regulation, and greater public opposition.

By contrast, transfers of water from existing sources can be used to more efficiently manage risk posed by rising demand, allowing water to be acquired on more of an as-needed basis. When developing transfer agreements, however, questions of timing, quantity, and type of transfers must be settled if transfers are to be effectively employed. Regional differences in water law, the nature of the available resources and the degree of hydrologic variability further determine how transfers might be applied.

This research contributes to knowledge in three specific areas:

- (i) This work examines the manner in which different types of market-based transfers can be combined with firm capacity to form minimum expected cost “portfolios” of different transfer types (e.g., permanent rights, leases, options)

that meet defined reliability and/or cost variability constraints. In doing so, a Monte Carlo simulation is paired with the “implicit filtering” optimization routine, designed to optimize portfolios despite the sampling error, or “noise”, inherent in searching for an optimal expected value.

- (ii) The second phase of research applies a modified technique (control variate) to reduce the level of noise inherent in the simulation, thereby improving the efficiency and accuracy of the optimization approach. This method is applied to the study region as the simulation is expanded from a one-year to a 10-year model, and results in a significant reduction in computational burden (as much as 50%).
- (iii) A technique is developed to generate synthetic streamflow time series in a manner that reproduces autocorrelation in the historic record. This method is used to develop streamflow records representative of future climate scenarios, which are then used as inputs for a model that assesses different risk-based transfer agreements within the Research Triangle region of North Carolina. Results demonstrate that even minor changes in expected streamflows can significantly impact transfer activity and costs.

To Emily, my loving wife

Table of contents

List of figures.....	vii
List of tables.....	ix
Chapter 1: Introduction.....	1
Chapter 2: Optimization of Water Supply Portfolios in the Lower Rio Grande Valley.....	6
2.1 Introduction.....	6
2.2 Methodology.....	11
2.2.1 Hydrologic-Market Simulation.....	13
2.2.2 Optimization Framework.....	21
2.2.3 Control Variate Method.....	26
2.2.4 Application of CV Method.....	28
2.2.5 Expansion to Multi-Year Model.....	31
2.2.6 Application of CV Method to Multi-Year Simulation.....	32
2.2.7 Study Region.....	35
2.3 Results.....	45
2.3.1 Developing Minimum Cost Portfolios.....	45
2.3.2 Application of CV Method to Single-Year Model.....	55
2.3.3 Multi-Year Portfolio Scenarios.....	60
2.3.4 Long-Term Option Contracts.....	64
2.4 Conclusions.....	67

Chapter 3: Improved Method for Developing Alternative Hydro-climate

Scenarios and their Effects on Water Transfer Agreements.....	69
3.1 Introduction.....	69
3.2 Methodology.....	74
3.2.1 Study Region.....	74
3.2.2 Managing Transfers.....	77
3.2.3 Autoregressive Models and Autocorrelation.....	79
3.2.4 The Autocorrelated Bootstrap.....	81
3.2.5 Application of the Model to the Study Region.....	91
3.3 Results.....	94
3.4 Conclusions.....	102
Chapter 4: Conclusions.....	104
References.....	108

LIST OF FIGURES

2.1 Expected cost landscape.....	22
2.2 Schematic of optimization algorithm, model, and pilot study.....	30
2.3 Map of Lower Rio Grande Region.....	36
2.4 Minimum cost portfolios for different strategies.....	46
2.5 Variations on strategy C2.....	53
2.6 Landscapes with and without CV method.....	56
2.7 Average optimized results of the one-year model.....	57
2.8 Optimized one-year portfolio results.....	59
2.9 Optimized 10-year portfolio results.....	60
2.10 Average optimized results of the 10-year model.....	62
2.11 Map of variance reduction.....	64
2.12 Expected option activity over 10 years.....	65
3.1 Risk chart for OWASA in 2025.....	82
3.2 How Y is converted to Y'	89
3.3 Partial autocorrelation function.....	90
3.4 The 1-lag correlation of historic OWASA inflows.....	91
3.5 Autocorrelation function.....	92
3.6 Histograms of Standardized Inflow Index.....	93
3.7 Histogram of SII values for high flow events.....	95
3.8 Expected annual volume of water transferred to Durham.....	98

3.9 CDFs of the lowest observed annual storage for OWASA.....	99
3.10 CDFs of the lowest observed annual storage for Durham.....	101

LIST OF TABLES

2.1 Summary of control variates used in multi-year scenarios.....	34
2.2 Simulation data summary.....	38
2.3 Minimum cost portfolios.....	47
2.4 Strategy B and C2 portfolios.....	52
3.1 Total costs of transfers.....	103

Chapter 1: Introduction

Rising water demand and concerns over scarcity have driven an increasing number of regions to explore market-based approaches to water resource management (Anderson and Hill 1997; Easter, Rosegrant et al. 1998; National Research Council 2001). Nonetheless, most water markets remain relatively unsophisticated, with transactions involving only permanent transfers of water rights. While a number of studies have shown that permanent transfers encourage long-term allocation efficiency (Howe, Schurmeier et al. 1986; Young 1986; Saliba 1987; Chang and Griffin 1992; Griffin and Boadu 1992; Colby, Crandall et al. 1993; Hearne and Easter 1997; Howe and Goemans 2003; Brookshire, Colby et al. 2004; Nieuwoudt and Armitage 2004), such transfers provide a less effective means of managing short-term scarcity. Rising demand in many regions has increased the level of economic and social disruption brought about by seasonal droughts, consequently some markets are beginning to support a more sophisticated menu of temporary transfers (Howitt 1998). In response, some researchers have investigated the potential efficiency gains associated with “spot market” leases (Vaux and Howitt 1984; Smith and Marin 1993; Characklis, Griffin et al. 1999) and options (Hamilton, Whittlesey et al. 1989; Michelsen and Young 1993; Watters 1995; Jercich 1997; Howitt 1998; McCarl, Dillon et al. 1999; Villinski 2004).

Spot market leasing generally involves the immediate transfer of “wet” water, with the lease price subject to considerable variability based on supply and demand conditions. A typical option agreement involves an initial payment that guarantees the

purchaser the right to lease water at a later date at an agreed upon “exercise” price. The certainty inherent in the exercise price can make options an attractive hedge against spot market price volatility, while providing the additional advantage of postponing transfer decisions (and full payment) until better information is available. Both leases and options improve market flexibility relative to permanent transfers alone, allowing water users to more rapidly adapt to changing conditions while meeting their reliability goals with a reduced volume of “firm” capacity. As leases and options have become more widely available, there has been increased interest in how water users might coordinate the use of these instruments to achieve the dual objectives of maintaining water supply reliability and lowering supply costs.

Riparian law may not extend the property rights necessary to permit water markets, but generally, water utilities that withdraw and treat water may sell that treated water to another utility (with limitations). However, these sorts of transfers do require some type of infrastructure in place (e.g., pipelines, interconnections between distribution systems), generally limiting transfer opportunities more than the raw water transfers in the western U.S. Despite the more limited opportunities, transfers in the eastern U.S. can often meet reliability objectives at lower costs than expansions of firm capacity, though these transfers tend to be established on a contractual basis.

Regardless of the environment, both market- and contract-based transfers require active management of when and how much water to transfer. In the two study regions in this work, anticipatory, risk-based decision rules are used to manage transfers. The decision rules are designed differently for each study region, recognizing the different transfer environments and limitations of the particular study region. Both decision rules

are predicated on the idea that transfers should be initiated prior to water supply levels reaching a crisis point, and both express some explicit or implicit concept of risk, the likelihood of water supply levels reaching that crisis point.

The models used in this work involve Monte Carlo techniques, an important choice given the nature of how the transfers are being used. Transfers are not being used as a primary, everyday source of water supply. Rather, they are used as supply augmentation during droughts, and as such represent infrequent events with a wide variability in their frequency and quantity of consumption. Performing hundreds, or even thousands, of realizations (i.e., repetitions) of the simulation provides greater precision in the estimates of expected costs and outcomes. Ultimately, this work provides a range of tools that will enable utilities to confidently assess the viability, costs, and effectiveness of potential transfer strategies.

Chapter 2 combines two published papers into a single, integrated work. Both papers address the optimization of a “portfolio” of water supply assets of a utility in the Lower Rio Grande Valley in Texas. A Monte Carlo simulation is created to determine expected outcomes of a range of scenarios and portfolio strategies. Traditionally, Monte Carlo simulations have been of limited use in water resources planning. As a sampling technique, Monte Carlo simulations produce sampling errors, such that traditional, gradient-based optimization methods are of little use. Without the ability to optimize Monte Carlo simulations, their use has languished. In this work, the Monte Carlo simulation is paired with an optimization method that has been designed to optimize a simulation that exhibits the sort of sampling errors produced by Monte Carlo simulations. The first of the two papers provides a proof-of-concept, as this is the first example of

this sort of simulation-optimization pairing within the water resources community. The paper also demonstrates how risk-based decision rules can effectively govern the use of temporary transfers, such as leases and options. This paper ultimately shows how the composition of optimized portfolios can shift as the priorities of the utility change (e.g., the relative importance of minimal cost versus cost volatility versus reliability). The second paper within the chapter continues the work of the first paper, but expands the one-year simulation period of the first paper to a 10-year simulation, which is more in line with the planning periods of utilities. In order to counter the increased computational burden, the simulation is paired with a variance reduction technique. This is the first example of a variance reduction technique being used as an aid to optimization, not simply to improve the accuracy of a simulation. As a result, the use of this variance reduction technique roughly halves the computational burden required to produce an optimized portfolio with a given level of accuracy and precision. Moreover, the longer simulation period allows the exploration of how long-term option contracts can be adjusted to both accommodate growth in demand for water and minimize cost.

The third chapter in this work turns away from the optimization aspect of the previous chapter, instead focusing on improving the simulation to better predict the performance of temporary transfers. Specifically, this work improves the ability of a synthetic time series to replicate the autocorrelation of the historic record. Streamflows commonly contain a significant amount of autocorrelation within the time series, and replicating that autocorrelation is vital to replicating long-term high- and low-flow events. Moreover, other techniques that generate synthetic records containing autocorrelation do so based upon an analysis of the autocorrelation function, which

relates the autocorrelation contained within the historic time series as a whole. These other techniques neglect the fact that, particularly with streamflows, autocorrelation can fluctuate seasonally as hydrologic conditions change (e.g., evapotranspiration, soil moisture). The technique developed in this work, the autocorrelated bootstrap, not only reproduces the historic autocorrelation function, but also the seasonal fluctuations in the correlation structure. The manner in which data is treated as part of the autocorrelated bootstrap produces an opportunity to adjust seasonal averages and standard deviations to reflect future climate change scenarios. Thus, given improved climate change predictions, the autocorrelated bootstrap becomes a tool in which those predictions can be translated into a meaningful form for water utilities, that is, how the performance of water supplies is altered.

The autocorrelated bootstrap is applied to the Research Triangle region of North Carolina. The model simulates the use and performance of interruptible contracts governing the sale of treated water between urban users. The ability of the autocorrelated bootstrap to replicate the historic rate of long-term extreme flow events is proven. The simulation itself then shows that the expected rate of transfers is highly dependent upon future average inflows. For instance, if average annual inflows were to drop just 7%, expected transfer rates double.

In total, the work here presents a range of tools to aid in both the simulation and optimization water resources. Specifically, the use of temporary transfers is shown to be an efficient and effective mechanism for ensuring the reliability of a utility's water supply.

Chapter 2: Optimization of Water Supply Portfolios in the Lower Rio Grande Valley

(Note: This chapter combines the work published in two papers:

Kirsch, B. R., Characklis, G. W., Dillard, K. E. M. and C. T. Kelley (2009). "More Efficient Optimization of Long-term Water Supply Portfolios," *Water Resources Research*, 45, W03414.

Characklis, G. W., Kirsch, B. R., Ramsey, J., Dillard, K. E. M. and C. T. Kelley (2006). "Developing Portfolios of Water Supply Transfers", *Water Resources Research*, 42(5), W05403.)

2.1 INTRODUCTION

Several previous studies have used either linear or stochastic programming techniques to identify combinations of supply alternatives, including infrastructure, transfers and conservation, that minimize the expected costs of meeting urban water demand (Lund and Israel 1995; Wilchfort and Lund 1997; Watkins and McKinney 1999; Jenkins and Lund 2000). In general, these methods have involved some form of two stage model in which the first step involves a hydrologic simulation that is used to establish a discrete set of supply scenarios. This information is combined with price and usage data to develop least cost combinations of long-term (e.g., reservoirs) and short-term supply alternatives (e.g., leases, options), with results suggesting that the coordinated use of short-term transfers can reduce costs.

This work focuses solely on market-based transfers, but expands on earlier studies by employing a simulation-optimization approach that allows for the exploration of some

issues that have received less attention in previous work. In particular, this work describes portfolio development from a perspective of a utility manager seeking to minimize water supply costs. In many earlier studies a city's decision to acquire water via leases or options and its actual receipt of the water occur within a single time period. While these periods have often been long enough (3 to 6 months) that this is not unreasonable, such an approach assumes that the city buys and acquires water at exactly the time it is needed or, alternatively, that the city has perfect information regarding its future needs at the time it makes a purchase. Even in a market where transactions can be completed quickly, such a scenario is at odds with the risk averse nature of utilities who will generally seek to augment supply in advance of a shortfall (i.e. without perfect information). Toward that end, this work identifies anticipatory decision rules, using the ratio of expected supply to expected demand as the basis for determining when (and how much) to lease/exercise. These rules could provide a utility with a decisionmaking framework for arriving at a least cost solution using information as it becomes available throughout the year.

Uncertainty with respect to spot market prices will also be a primary concern when developing portfolios that include temporary transfers. In this work spot lease prices are represented as distributions (based on actual market data) and this information is used to price options in a risk neutral manner consistent with financial theory (Black and Scholes 1973; Hull 1999). The use of lease price distributions provides the additional benefit of allowing for the calculation of both expected cost and cost variability. This is a potentially important distinction because while minimizing expected supply costs is certainly important, it is likely that cost variability will also play a role in

decisions regarding a portfolio's suitability. Cost variability has been considered before, Watkins and McKinney (1997) describe a relatively elegant approach that incorporates consideration of both expected supply costs and their variance when identifying an optimal solution, however, their approach assumes a symmetrically distributed objective function (essentially cost). This may not be the case in many regions, a point made more significant given that the risk of high costs associated with asymmetric tails in the distribution may have a significant impact on decisionmaking.

The modeling approach employed consists of a hydrologic-market simulation embedded within a search-based optimization algorithm. This methodology is designed to identify the portfolio of rights, options and/or leases that minimizes expected costs while meeting constraints related to both supply reliability and (in some cases) cost variability. When minimizing the expected costs of water supply in a stochastic environment, computational burden can be a particular concern. In water supply problems, the expected cost surface near the optimum is often relatively flat and can be somewhat "noisy", increasing the likelihood that a search will become trapped in a local minimum. To combat these challenges, a different type of search technique ("implicit filtering") is used, one proven to be widely applicable for problems where the solution surface exhibits high-frequency, low-amplitude noise (Choi, Eslinger et al. 1999).

The computational burden of this type of simulation-optimization approach has traditionally been a deterrent to its broader use. The initial investigation made into portfolio optimization uses a short-term (i.e. one year) planning horizon as a proof-of-concept and to better differentiate the effects of various decision variables and initial conditions. However, water utilities' planning horizons are typically much longer. The

second phase of this research (i.e., the second paper) extends the planning horizon and develops techniques to reduce the now expanded computational burdens of the simulation-optimization.

Implicit filtering is used to navigate local minima caused by high-frequency, low-amplitude noise, and the ability to navigate this noise is dependent upon its magnitude. The magnitude of the noise is a product of the variance of the objective function (i.e. the expected portfolio cost), which can be reduced by increasing the number of realizations (i.e. simulation runs) on which each expected value is based but at the cost of an even greater computational burden. As a means of reducing this burden, a variance reduction technique known as the control variate (CV) method (Lavenberg and Welch 1981; Avramidis and Wilson 1996) is applied. The CV method can reduce the noise inherent in optimization surfaces based on expected values by using *a priori* knowledge of how random variations in a simulation's inputs affect its output, and using this information to reduce the variance in simulation output. The method has been used in other areas, most notably the pricing of stock options (Boyle 1977; Johnson and Shanno 1987; Broadie and Glasserman 1996), but always with a standalone Monte Carlo (MC) simulation. In this case, it is applied within a simulation-optimization framework with the intent of characterizing its potential for improving optimization efficiency. This is also the first application of this technique in a water resources context, so in addition to providing information on the degree of improved efficiency that might be expected, it should also provide some general insights into what factors are most important in improving computational efficiency for this class of problems.

The expansion of this concept to multi-year scenarios extends the planning horizon to a level of more interest to utilities, provides insights into how portfolio composition might change over time, and produces results that are less dependent upon initial conditions. It also allows for an investigation of several types of long-term option contracts, those which allow the purchaser with year-to-year flexibility, while still providing the long-term contractual security that frees the utility from the cost and inconvenience of annual renegotiation. Results suggest that the combination of implicit filtering and the control variate method is capable of significantly improving the efficiency of simulation-based optimization, a finding that could be applied in a broad range of water resource contexts. Similarly, results related to the composition of long-term water supply portfolios, including multi-year option contracts, may provide insights valuable in the formulation of water supply strategies.

This simulation-optimization approach is applied to the Lower Rio Grande Valley, a region that supports an active water market (Griffin and Characklis 2002). The availability of hydrologic information and 10 years of spot lease price data make this region well suited for an exploration of water supply portfolio development. The region also exhibits characteristics typical of many water scarce western regions, including rapidly growing municipal demand and a large agricultural sector. Results should provide general insights into the role that options and leases can play in lowering the cost of meeting water supply reliability goals. While this work represents firm supply capacity as water rights, a similar approach could be used to develop portfolios integrating options and/or leases with any form of hard supply infrastructure.

2.2 METHODS

An approach is developed to identify a minimum cost portfolio of rights and transfers that meets one city's water demand with a specified reliability over a period of 12 months. The regional water supply is provided via a reservoir, with water allocated to users through a system of rights. Water can be obtained via the following three mechanisms:

The first is permanent rights. These entitle the holder to a pro rata share of reservoir inflows (after correcting for losses), such that a city owning 5% of regional rights is allocated 5% of inflows. Allocations are made at the end of each month, and the water can be used in any subsequent month. Permanent rights are transferable, but regulatory approval takes time, so the city's volume of permanent rights is assumed constant throughout the year. Their price (p_R) is represented as an annualized cost based on purchase price.

The second mechanism is spot market leases. Lease transactions can be completed at the end of each month, and leased water may then be used in any subsequent month. Leasing transactions receive less regulatory scrutiny as they involve only a temporary transfer and so may be completed quickly (i.e., within a few days). Spot lease prices in each month t are linked to reservoir levels and described as random variables (\hat{p}_{L_t}).

The third mechanism is option contracts. Option contracts provide the right to lease water at a later date and an agreed upon price. Options can be purchased just before the beginning of the year and "exercised" on a single call date (i.e., a European call option) that corresponds to the last day of a specified month (t_X). Once an option has

been exercised, the leased water can be used in any subsequent month. Options not exercised on the call date lapse and have no further value. Option prices (p_o) and exercise prices (p_x) are based on the distribution of spot lease prices (\hat{p}_{L_t}) in the exercise month.

Options are priced using a “risk-neutral” approach in which it is assumed impossible to make risk-free profits (Black and Scholes 1973). In other words, the expected value an option provides relative to a spot market lease does not exceed the option’s price (Hull 1999). The price of a European call option (p_o) is calculated by discounting the option’s expected value on the call date back to the point at which the option is purchased, with the expected value based on the difference between the exercise price and spot lease price, or zero, whichever is larger (expression in brackets), such that

$$p_o = e^{-rT} \cdot E \left[\max(\hat{p}_{L_t} - p_x, 0) \right] \quad [2.1]$$

where r is the discount rate (monthly) and T is the period between purchase and exercise dates (months).

The general approach to portfolio development first involves constructing a stochastic simulation that models the city’s responses to changing hydrologic and market conditions. The simulation is embedded within an optimization framework which, for any given set of initial conditions, identifies the portfolio of water market transfers that minimizes expected costs while meeting constraints related to reliability and cost variability. The regional context is the western United States, a setting where agricultural water use generally dominates and increasing water scarcity is driven by urban expansion. As such, there are several implicit assumptions. One is that the city is a relatively small player within the regional market and exercises no market power (i.e., it

is a price taker). In addition, because the vast majority of water is used for relatively low value irrigation, it is assumed that the city can always find sufficient water available within the market to accommodate a lease or exercise transaction. It is worth noting that while the assumptions related to the unlimited availability of spot market water and risk-neutral pricing provide a reasonable basis for this analysis, their use may have implications for results, and these will be discussed in later sections.

2.2.1 Hydrologic-Market Simulation

The simulation runs over a 12-month period, beginning on 31 December ($t = 0$), with the city holding some number of permanent water rights (N_{R_T}) and options (N_O). Initial conditions specify reservoir storage (R_0) and the amount of water the city has carried over from the previous year (N_{t_0}). In each of the following months, regional hydrologic conditions are simulated using data sets describing monthly reservoir inflow, outflow, and losses, with these conditions linked to both the city's water supply and the spot market price for water. This information is then combined with monthly distributions of the city's demand to make decisions regarding the purchase of leases and/or exercise of options. Multiple simulation runs for each set of initial conditions generate values for the expected annual cost of the city's portfolio, expressed as (random variables are denoted by the circumflex)

$$E[Annual\ Cost] = N_{R_T} p_R + N_O p_O + E[N_X] p_X + E\left[\sum_{t=0}^{11} N_{L_t} \hat{p}_{L_t}\right] \quad [2.2]$$

where,

N_{R_T} = total volume of permanent rights held by city (ac-ft);

N_o = volume of options purchased at the beginning of the year (ac-ft);

N_x = volume of exercised options (ac-ft);

N_{L_t} = volume of spot leases purchased at the end of each month (ac-ft).

Within the simulation, the following constraints apply:

$$N_x \leq N_o \quad = \text{the city cannot exercise more options than it buys in } t = 0; \quad [2.3]$$

$$\sum_{i=0}^{11} N_{r_i} \leq N_{R_T} \quad = \text{allocations of reservoir inflows to the city's permanent rights cannot exceed the number of rights that the city holds;} \quad [2.4]$$

$$R_{Max} \geq R_t \geq R_{Min} \quad = \text{Reservoir level must stay within specified bounds related to storage capacity } (R_{Max}) \text{ and minimum storage levels } (R_{Min}); \quad [2.5]$$

Non-negativity constraints also apply for all variables.

A series of variables are used to describe regional hydrologic conditions, including i_t = volume of reservoir inflows for each month t ; l_{R_t} = volume of reservoir losses for each month t ; and o_t = volume of reservoir outflows (including spillage) for each month t .

A water balance is maintained on the reservoir system throughout the simulation such that

$$R_t = R_{t-1} + i_t - o_t - l_{R_t}. \quad [2.6]$$

From the perspective of the individual city, total reservoir storage is less important than the volume of water available to the city itself, an amount largely determined by the city's initial supply (N_{r_0}) and its share of monthly reservoir inflows (N_{r_t}). Reservoir inflows available for allocation are calculated as the difference between monthly inflows and losses, multiplied by an instream loss factor (l_t), which accounts for losses incurred between the reservoir and user (which in this case is assessed prior to allocation). Inflows available for allocation to rights holders in each month (\hat{n}_t) are computed as

$$\hat{n}_t = (\hat{i}_t - \hat{l}_t) \bullet (1 - l_t) \quad [2.7]$$

These inflows are allocated on a pro rata basis such that the distribution of new monthly inflows accruing directly to the city (\hat{N}_{r_t}) is represented as

$$\hat{N}_{r_t} = \hat{n}_t \bullet \left(\frac{N_{R_T}}{\bar{N}_R} \right) \quad [2.8]$$

where,

\bar{N}_R = total volume of regional water rights.

The total volume of water available to the city in any month is assessed at the end of the preceding month, and the method of calculation changes depending on whether it is before or after the exercise month (t_x). In months prior to t_x , the supply available to the city in the next month (S_{t+1}) includes cumulative inflows and purchased leases, less water usage such that

$$S_{t+1} = \sum_{i=0}^t N_{r_i} + \sum_{i=0}^{t-1} N_{L_i} - \sum_{i=1}^t u_i, \quad \text{for } t = 0, 1, 2 \dots t_x - 1. \quad [2.9]$$

where,

u_t = city's usage in month t .

In subsequent months, the available supply also includes exercised options, such that

$$S_{t+1} = \sum_{i=0}^t N_{r_i} + \sum_{i=0}^{t-1} N_{L_i} - \sum_{i=1}^t u_i + N_X \quad \text{for } t = t_X, t_X+1 \dots 11. \quad [2.10]$$

The decision of whether or not to purchase leases is the last step in each month, and the decision is based on the city's available supply, specified by [2.9] or [2.10] (neither of which include consideration of leases purchased in month t). The leasing decision involves consideration of both the city's available supply and the volume of monthly inflows it expects to have allocated to it over the remainder of the year (calculated on the basis of historical records). These two values are summed to yield the city's expected water supply ($S_{E_{t+1}}$) over the remainder of the year, such that

$$S_{E_{t+1}} = S_{t+1} + \sum_{i=t+1}^{11} E[\hat{N}_{r_i}] \quad \text{for } t = t_X, t_X+1 \dots 10. \quad [2.11]$$

where \hat{N}_{r_i} = distribution of inflows allocated to the city in each month t .

November ($t = 11$) inflows are considered when calculating the available supply for December, but December inflows are allocated to the following year. Therefore December's available supply and expected supply are equal (i.e., $S_{E_{t+1}} = S_{t+1}$).

Once the city's expected water supply has been calculated, the decision is made to purchase leases and/or exercise options. This is a two-part decision in which the first step involves determining whether or not to acquire water and the second involves deciding how much. Both decisions are based on the ratio of expected supply to expected demand,

with the decision to acquire made by comparing this ratio against a specified threshold value (α), such that if

$$\frac{S_{E_{t+1}}}{\sum_{i=t+1}^{12} E[\hat{d}_i]} \leq \alpha, \text{ then, the city will acquire water, for } t = 0, 1, 2 \dots 11 \quad [2.12]$$

where,

\hat{d}_t = distribution of the city's water demand during each month t .

The question of how much to lease and/or exercise is made by comparing the ratio of expected supply to expected demand with a second specified threshold value (β). This leads to leases (N_{L_t}) being purchased and/or options (N_X) exercised until

$$\frac{(N_{L_t} + N_X) + S_{E_{t+1}}}{\sum_{i=t+1}^{12} E[\hat{d}_i]} = \beta, \text{ for } t = 0, 1, 2 \dots 11 \quad [2.13]$$

In all months except t_x , $N_X = 0$ and the volume of leases purchased can be represented as,

$$N_{L_t} = \beta \left(\sum_{i=t+1}^{12} E[\hat{d}_i] \right) - S_{E_{t+1}}, \text{ for } t \neq t_x. \quad [2.14]$$

During t_x , the decision process is modified such that exercising options is considered before purchasing leases. Under these conditions, the first step is to compare the exercise price (p_X) with the current spot lease price (p_{L_t}). If the lease price is less than the exercise price, the city will simply lease the volume defined in [2.14]. If, however, the exercise

price is less than the lease price, the city will exercise options, with the volume to be exercised expressed as follows:

$$\text{if } \beta \left(\sum_{i=t+1}^{12} E[\hat{d}_i] \right) - S_{E_{t+1}} \leq N_O, \text{ then } N_X = \beta \left(\sum_{i=t+1}^{12} E[\hat{d}_i] \right) - S_{E_{t+1}}, \text{ otherwise } N_X = N_O. \quad [2.15]$$

In the case of the latter scenario, where options alone are insufficient to satisfy [13], the city will acquire additional water via leasing, such that

$$N_{L_t} = \beta \left(\sum_{i=t+1}^{12} E[\hat{d}_i] \right) - S_{E_{t+1}} - N_X, \text{ for } t = t_X. \quad [2.16]$$

Different α and β variables can be specified for individual seasons or even individual months. In the example described later, two different parameter pairs are established, one (α_1/β_1) for the period running up to the month before options can be exercised ($t_0 \rightarrow t_X - 1$) and another (α_2/β_2) for the remainder of the year. Expected supply [2.11] is similarly partitioned, such that it is calculated relative to t_X in months leading up to t_X , and calculated relative to the end of the year in all subsequent months. Optimal values for α and β , those that lead to a minimum expected cost portfolio that meets reliability constraints, are determined as part of the optimization routine (see next section).

The choice to link decision rules to the ratio of expected supply to expected demand was based on the ability to use this value in determining both when and how much water to acquire. Alternative decision rules could have been based on the probability of shortfall, or perhaps even linked to a threshold value for the expected benefits loss that would accrue as a result of a shortfall. These types of rules may be expressed in terms more intuitive to utility personnel and/or planners (and might be

explored in future work), but their use would have necessitated additional calculations to answer both the “when” and “how much” questions.

Water is acquired just before the monthly counter changes (i.e., month $t + 1$ becomes month t), correspondingly $S_{t+1} \rightarrow S_t$, which is then represented as

$$S_t = \sum_{i=0}^{t-1} N_{r_i} + \sum_{i=0}^{t-1} N_{L_i} - \sum_{i=1}^{t-1} u_i, \quad \text{for } t = 1, 2 \dots t_X, \quad [2.17]$$

or

$$S_t = \sum_{i=0}^{t-1} N_{r_i} + \sum_{i=0}^{t-1} N_{L_i} - \sum_{i=1}^{t-1} u_i + N_X \quad \text{for } t = t_X+1, t_X+2 \dots 12. \quad [2.18]$$

Available supply (S_t) is compared with a demand value (d_t) obtained by either randomly sampling a monthly distribution or selecting from a monthly time series. If available supply is sufficient to meet this demand (i.e., $S_t \geq d_t$), then demand equals usage ($u_t = d_t$). If available supply is insufficient, then $u_t = S_t$, leaving a shortfall of $d_t - S_t$ and a “failure” is recorded for that month. A distinction is made between a “failure” and a “critical failure” ($S_t/d_t \leq 0.6$) in order to recognize differences in severity and the measures that would be required to compensate for the shortfall. A running tally of both failures and critical failures is maintained throughout the simulation.

Once available supply and demand have been compared, the process of evaluating new allocations and lease/exercise decisions repeats monthly through the end of the year. Each annual run within this probabilistic framework represents one realization of the cost and reliability of a portfolio defined by selected values for the initial conditions (R_0, N_{r_0}) and decision variables ($N_R, N_O, \alpha_1, \beta_1, \alpha_2, \beta_2$). Multiple runs are made to determine a portfolio’s expected cost (equation [1.2]) and expected reliability, with the latter defined as

$$E[r_f] = 1 - \left(\frac{\text{failures}}{12 \bullet \text{Years}} \right) \quad [2.19]$$

where,

r_f = monthly reliability against a failure (i.e. $S_t < d_t$);

Years = number of simulated years (i.e. annual runs).

A reasonable span of monthly reliabilities might range from 0.995 (i.e., one failure every 16.7 years) to 0.98 (one failure every 4.2 years). A similar factor (r_{cf}) is used to measure the expected reliability relative to critical failures.

Multiple annual runs also allow for evaluation of the probability of very high annual costs. Within the electricity and natural gas industries, a common metric used to describe the risk of high costs is the “contingent value at risk” (CVAR). Given a distribution of annual costs, the CVAR represents the mean of the annual costs falling above the 95th percentile. Something akin to the CVAR is likely to play a role in utility decisions, and this metric is used here.

The quantity of water remaining in the city’s possession at year’s end is also tracked. This remaining water is not assigned any value, a shortcoming that could raise concerns that a portfolio developed within this annual framework may not bear much resemblance to the type of portfolio that would minimize costs over a longer time horizon. For instance, a portfolio that consistently left the city with very little water at the end of the year could result in very high supply costs the following year (this does not actually tend to be the case, however). While the development of long-term portfolios is beyond the scope of this work, these issues will receive some attention in the results section.

The methodology described above involves a supply strategy that includes rights, options, and leases (strategy C); however, it is easily modified to explore alternative strategies that include permanent rights alone (strategy A) and permanent rights and options (strategy B). In the case of a city relying on strategy A, the number of rights (N_R) becomes the only decision variable. With respect to strategy B, the number of decision variables increases to four ($N_R, N_O, \alpha_2, \beta_2$) and the decision framework for acquiring water (i.e., equations [2.12], [2.13], and [2.15]) is similar to that described above, except that the city acquires additional water via options alone, and only in the exercise month. Strategy C involves six decision variables (α_1, β_1 are added) and the entire monthly decision framework described above.

2.2.2 Optimization Framework

The simulation is linked to a search algorithm that identifies optimal values for the decision variables based on the following formulation (for Strategy C),

$$\underset{N_R, N_O, \alpha_1, \beta_1, \alpha_2, \beta_2}{\text{Minimize}} \quad Z = E[\text{Annual Cost}] \quad [2.20]$$

Such that:

$$E[r_f] \geq \text{monthly reliability threshold}, \in [0,1]; \quad [2.21]$$

$$E[r_{cf}] \geq \text{monthly critical reliability threshold}, \in [0,1]. \quad [2.22]$$

Later results also incorporate an additional constraint limiting cost variability, such that

$$\frac{CVAR}{E[\text{Annual Cost}]} \leq \text{cost risk threshold}, \in [1,\infty). \quad [2.23]$$

Figure 2.1 illustrates a section of the optimization landscape describing expected cost as a function of the number of permanent rights and options ($\alpha_1, \beta_1, \alpha_2, \beta_2$ held constant). While the surface is relatively smooth when the volume of leases and exercised options is small (i.e., when a portfolio is mostly rights), as the volume of leases and exercised options increases so does the “noise”. This can be problematic for many gradient-based search algorithms as they can become trapped in local minima. The amplitude of the noise can be reduced by increasing the number of simulated years, but this comes at a price in terms of computational burden.




Figure 2.1: Expected cost landscape (constant values for α_1/β_1 and α_2/β_2)

Implicit filtering is a finite difference search method in which the difference increment (i.e., the size of the finite difference stencil) is varied as the optimization

progresses (Kelley 1999). In this way, local minima which are artifacts of low-amplitude noise do not trap the iteration, and the noise is “implicitly filtered” out. This is in contrast to methods which explicitly try to filter out high-frequency components of the objective function (Kostrowicki and Piela 1991; More and Wu 1997); such methods are designed for problems with high-amplitude high-frequency terms and should be thought of as global optimization algorithms. Implicit filtering is not a global optimization method, and is designed to efficiently solve problems, such as those presented in this paper, which have noisy but not violently oscillatory optimization landscapes (see Figure 2.1). Methods such as steepest descent, which are based on gradients, can be trapped in the small-scale local minima that noisy surfaces exhibit, and may fail if this results in an optimization surface that is not differentiable. In this problem, as in many others, the noise results from using an expected value (cost) as the objective function. The frequency and amplitude of the noise increases with greater use of leases and exercised options (probabilistic variables) and decreases with the number of simulated years used to generate an expected cost estimate of each portfolio. While an infinite number of simulations for each portfolio would generate a perfectly smooth optimization surface (which could be optimized using some form of steepest descent approach), implicit filtering allows for efficient optimization of the problem by allowing the search to progress while reducing the number of simulated years required to generate expected cost values during each iteration.

Implicit filtering uses the finite difference gradient (as described by the difference between points on the finite difference stencil) to compute a search direction for descent. Unlike the classical steepest descent method, in which the negative gradient (or an

approximation of the negative gradient) is used, implicit filtering uses a quasi-Newton model of the Hessian to scale the gradient, thereby accelerating convergence in the terminal phase of the iteration. The theory for implicit filtering (Stoneking, Bilbro et al. 1992; Kelley 1999) and related algorithms (Torczon 1997; Kelley 1999; Audet and Dennis 2003) explains how such methods overcome low-amplitude noise and also gives insight into the limitations of these methods. In particular, there is no guarantee that a global minimum will be found. While implicit filtering cannot ensure convergence to a global minimum (this can only be proven for methods that undertake exhaustive efforts to asymptotically sample a dense subset of the design space), there is a rich literature describing the convergence of this class of methods, generally distinguished by the “polling” of stencil points throughout an iteration (Torczon 1997; Kelley 1999; Audet and Dennis 2003). This body of work demonstrates that for problems involving a smooth objective function and inequality constraints, any limit point of an iteration satisfies the first-order necessary conditions for optimality, which is the typical conclusion in convergence theorems for iterative methods for optimization. These results have also been generalized to both nonsmooth (Audet and Dennis 2003; Finkel and Kelley 2004) and noisy problems (Stoneking, Bilbro et al. 1992; Choi, Eslinger et al. 1999).

In this application, the implementation code, implicit filtering for constrained optimization (IFFCO), uses the difference gradient stencil for more than computation of the gradient (Choi, Eslinger et al. 1999). The gradient-based optimization is augmented with a coordinate search using the stencil points. If the result of the coordinate search is better than the result from the descent method, IFFCO accepts the coordinate search result. The coordinate search is also used in one of the termination tests for optimization

(for details, see Choi et al. [1999] and Kelley [1999]). IFFCO handles constraints in two ways. Simple bound constraints on variables (e.g., $N_o \geq 0$) are enforced at each iteration by setting variables that exceed the bounds to the value of the nearest bound. Indirect constraints (e.g., reliability) are handled by assigning slightly higher values to the objective function of points where the constraint is violated. These failed points are always at the edges of the stencil, and they act to steer the search away from the infeasible region. IFFCO's combination of stencil-based sampling and gradient-based optimization is most effective when the function to be minimized is a smooth surface with low-amplitude perturbations. Such problems are common in a number of applications, and while implicit filtering has not been applied to water resource management problems, it has been successfully employed in some related settings, including the design of groundwater remediation systems (Batterman, Gablonsky et al. 2002; Fowler, Kelley et al. 2004).

The simulation-optimization procedure includes 10,000 annual simulation runs for each set of decision variables, generating values for expected costs, reliability, critical reliability, and the CVAR which are generally reproducible to three significant figures. These parameters, as well as the α and β values, are passed to IFFCO which then guides the search of the optimization landscape. A search duration of 50 calls to the function (i.e., simulation) per decision variable was generally found to provide a resolution with respect to the expected cost and portfolio composition that corresponded to less than 1% and 200 ac ft, respectively. In some cases, 50 calls were insufficient to reach this resolution, and in these instances the solution from the first 50 calls (or a close

approximation) was used as a starting point and the process repeated until changes in the solution were within these tolerances.

2.2.3 Control Variate Method

Despite the ability of implicit filtering to navigate noisy surfaces, the noise produced by an objective function ($f(\mathbf{x})$) is detrimental to the precision and accuracy of the optimized solution. This noise can be quantified using the standard error of the mean (*s.e.*), defined as,

$$s.e. = \frac{\sigma}{\sqrt{n}}, \quad [2.24]$$

where σ is the standard deviation of $f(\mathbf{x})$ and n the number of realizations. As the value of σ is intrinsic to $f(\mathbf{x})$, controlling the noise associated with a MC simulation typically means controlling the number of realizations performed. However, the square root in the denominator means that increasing n has decreasing marginal returns in reducing standard error. An alternative approach to reducing noise would be to reduce the variance of the function, something that can be achieved through application of the control variate method.

The control variate (CV) method is a variance reduction technique that utilizes knowledge of how variation of stochastic input variables affects the value of simulation output, in this case, the objective function $f(\mathbf{x})$. For example, a city may wish to estimate the expected cost of transfers through MC simulation. If there exists a known correlation between the cost of water transferred (simulation output) and the volume of water stored in the city's reservoir (stochastic input variable), the reservoir storage volume may be used as a control variate. Given a known mean reservoir storage volume, deviations from

its mean can be used to account for deviations from the estimated mean cost of transfers, thereby reducing the variance in the transfer costs and improving the precision of cost estimation. While a more in-depth discussion can be found in the literature (Lavenberg and Welch 1981; Avramidis and Wilson 1996), a brief description is offered here.

If Z is a random input variable that is sufficiently correlated to model output, $f(\mathbf{x})$, θ can be defined as the variance-reduced value of $f(\mathbf{x})$ such that

$$\theta = f(\mathbf{x}) + c \cdot (Z - E[Z]), \quad [2.25]$$

where c is a scaling factor and Z is the control variate. Taking the expected value of both sides of [2.25] produces

$$E[\theta] = E[f(\mathbf{x})], \quad [2.26]$$

such that θ becomes an unbiased estimator of $f(\mathbf{x})$ when c is any real number. If the variance of both sides of [2.25] is calculated, the following is obtained:

$$\text{Var}(\theta) = \text{Var}(f(\mathbf{x})) - 2c\text{Cov}(f(\mathbf{x}), Z) + c^2\text{Var}(Z). \quad [2.27]$$

It can be shown that if

$$2c\text{Cov}(f(\mathbf{x}), Z) > c^2\text{Var}(Z), \quad [2.28]$$

then θ has lower variance than $f(\mathbf{x})$. Further, it can be shown that minimum variance occurs at

$$c^* = \text{Cov}(f(\mathbf{x}), Z) / \text{Var}(Z). \quad [2.29]$$

The reduction in variance then can be predicted with

$$\text{Var}(\theta) = (1 - \rho^2)\text{Var}(f(\mathbf{x})), \quad [2.30]$$

where ρ is the correlation coefficient between $f(\mathbf{x})$ and Z .

The control variate method can be extended to accommodate multiple control variates (Z_1, Z_2, \dots, Z_j) , through the expansion of [2.25], such that

$$\theta = f(\mathbf{x}) + c_1 \cdot (Z_1 - E[Z_1]) + c_2 \cdot (Z_2 - E[Z_2]) + \dots + c_j \cdot (Z_j - E[Z_j]). \quad [2.31]$$

Similarly, the variance of θ is minimized through the choice of optimal values for c_1, c_2, \dots, c_j .

For the purpose of readability, references to the output variable $f(\mathbf{x})$ in this discussion will be replaced with expected cost ($Cost$), the output variable (or objective function) specific to this work. Likewise, the variance reduced output variable produced by the CV method, θ , will be replaced with $Cost_{var}$, such that [2.25] could be rewritten as:

$$Cost_{var} = Cost + c \cdot (Z - E[Z]). \quad [2.32]$$

2.2.4 Application of CV Method

Selection of appropriate control variates is guided by the modeler's understanding of sources of variability in the objective function. In this case, the objective function is expected portfolio cost, $Cost$, and the source of the variability in the portfolio cost arises from the purchase of leases and exercise of options. More specifically, the variability can be identified as arising from variability in both the price and the quantity of transfers acquired, both of which are linked to variability in reservoir inflows and water demand.

Two control variates are used in the one-year simulation. The majority of transfers occur at two decision points, the beginning of the year (t_0) and in May, the option exercise month (t_5). The lease price at t_0 is a function of a known distribution with a known expected value, obtained from water market lease price data (Watermaster's Office 2004). Because each individual realization begins at t_0 with the initial conditions known, the quantity of leases purchased are unchanged from realization to realization

(unless initial conditions at t_0 are changed). Thus, controlling for the variability in the lease price at t_0 accounts for all the cost variability that arises from leases purchased then, and the lease price at t_0 is designated as the first control variate, Z_{PL} .

The second control variate accounts for portfolio cost variability arising from variability in the quantity of transfers acquired in t_5 . Within the simulation, both the monthly rate of new reservoir inflows allocated to the city ($N_{R,i}$) and the city's monthly water demand (D_i) have known expected values, and the difference between the two is the monthly *net supply*. The second control variate, Z_{NS} , is thus defined as the net supply from the beginning of the year (t_0) to t_4 , the month prior to the option exercise month, such that

$$Z_{NS} = \sum_{i=0}^{t_4} N_{R,i} - D_i . \quad [2.33]$$

Therefore, below average values of Z_{NS} indicate above average lease purchasing or option exercising activity in t_4 . Incorporating [2.33] into [2.32], the variance-reduced cost estimate for the one-year model ($Cost_{var}$) can be represented as

$$Cost_{var} = Cost + c_1 \cdot (Z_{PL} - E[Z_{PL}]) + c_2 \cdot (Z_{NS} - E[Z_{NS}]) . \quad [2.34]$$

The optimal values of c (c^*) in [34] are not known *a priori* and will change with different decision variables and initial conditions. Therefore, values for c_1^* and c_2^* are estimated for each new set of conditions using a pilot study, involving a very limited number of realizations that produce correlations between the control variates and $Cost$. Figure 2.2 illustrates how the optimization algorithm, the model, and the pilot study relate

to each other. Without the CV method, the optimizer queries the model with an \mathbf{x} , a vector describing all six decision variables, and the model returns $Cost$. With the CV method, the primary simulation (within the ‘Main Model’ in Fig. 2.2) immediately passes \mathbf{x} to the pilot study, which performs a small number of realizations, calculates the c^* values and returns them to the main model. The main model then performs the primary simulation and applies [2.34] with the calculated c^* values before returning the variance-reduced cost estimate ($Cost_{var}$) to the optimizer. While the pilot study represents a computational investment, it is generally a small investment, and one that pays off in a decrease in the total number of realizations that must be performed.

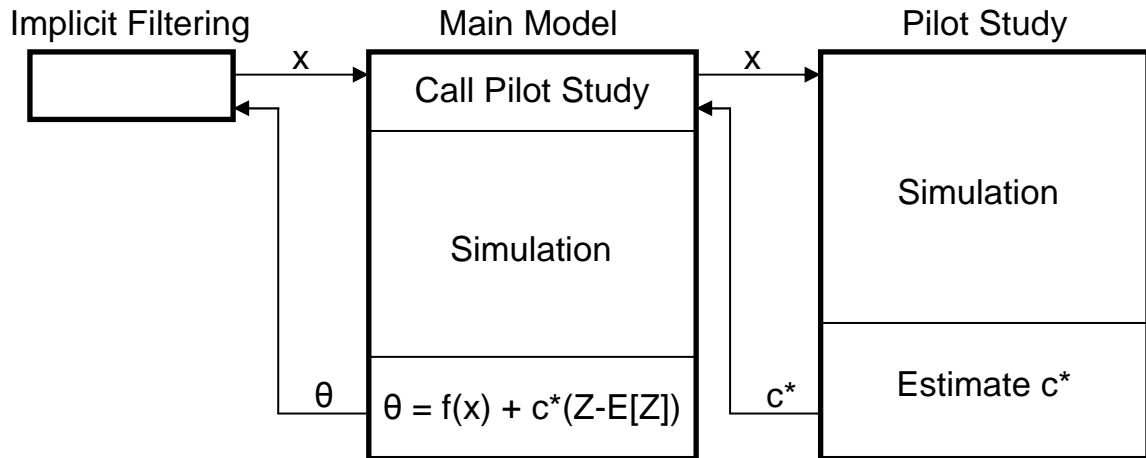


Figure 2.2. Schematic of optimization algorithm, model, and pilot study.

2.2.5 Expansion to Multi-Year Model

A multi-year simulation allows temporary transfers, particularly option contracts, to be evaluated on a time-scale that may be more useful for water supply planners. While the expanded model can accommodate any number of years, a 10-year planning horizon is used here.

From an optimization standpoint, the greatest change made to the simulation is reflected in the objective function, which is now represented in a multi-year form

$$\sum_{k=1}^{10} Cost_k = \sum_{k=1}^{10} [N_{R_T,k} p_R + N_O p_{O,k} + E[N_{X,k}] p_X] + E \left[\sum_{k=1}^{10} \sum_{i=0}^{11} N_{L_i,k} \hat{p}_{L_i} \right] \quad [2.35]$$

where k is the simulation year.

The operation of the hydrologic portion of the simulation is similar to the single-year simulation with some exceptions. The multi-year simulation is set up to account for annual growth (r) by multiplying each demand value by the term $(1 + r)^{k-1}$. In addition, the reliability constraint is modified to accommodate multiple simulation years, such that the reliability for each year within the simulation period is required to meet a minimum value. The cost variability constraint in the multi-year scenario is also re-defined such that the average annual ratio of CVAR-to-expected cost must be less than a specified value.

In addition to exploring how a longer planning horizon alters optimal portfolio composition, the multi-year model presents opportunities for examining long-term option contracts, all of which operate on as a series of one-year contracts with agreed upon

provisions to accommodate growth. For example, an annual growth factor can be attached to the volume of options purchased each year. This growth factor can be calculated to increase in accordance with rising demand. In this case, model input is changed such that N_O refers to just the number of options purchased in the first year, with subsequent years' option purchases defined as

$$N_{O,k+1} = D_{E,k+1} - D_{E,k} + N_{O,k} \quad . \quad [2.36]$$

One last point involves the way in which costs are presented. The budgeting cycle for utilities is typically annual, driven by a desire to recover costs. Consequently, costs are presented in annual terms. However, some recognition of the likelihood that some assets' costs will accrue over multiple time periods is appropriate. The cost of permanent rights is annualized (over 20 years using a 6 percent discount rate) as would be consistent with purchases funded by municipal bonds. The costs of leases and options are incurred only in the year in which they are bought or exercised (lease and option agreements expire after one year). The total portfolio cost (the sum of 10 annual costs) is presented in undiscounted terms as a way of representing the growth in annual costs over time. While these could be provided in discounted terms, this would not be consistent (outside of the annualized permanent rights costs) with the way in which water utilities will evaluate their water supply alternatives.

2.2.6 Application of CV Method to Multi-Year Simulation

The expansion of the model to a multi-year simulation requires several changes in how the CV method is applied. In the single-year simulation, the CV method is applied in order to reduce the variance of the objective function, $\text{var}(\text{Cost})$, whereas in the multi-

year simulation the goal is the reduction of $\text{var}(\sum_{k=1}^{10} Cost_k)$. This is accomplished by applying the CV method to each simulation year separately and calculating a reduced-variance cost ($Cost_{\text{var},k}$) for each $Cost_k$. Thus, the overall variance of each $Cost_{\text{var},k}$ must be sufficiently reduced such that

$$\text{var}(\sum_{k=1}^{10} Cost_{\text{var},k}) < \text{var}(\sum_{k=1}^{10} Cost_k). \quad [2.37]$$

In the first year of the multi-year simulation, the control variates used to calculate $Cost_{\text{var},1}$, the lease price in t_0 (Z_{p_L}) and the net supply of new water allocations from t_0 to t_4 (Z_{NS}), remain identical to those used in the single-year simulation. The lease price distribution, however, is dependent upon the reservoir level, but the expected reservoir level at t_0 of year $k + 1$ is dependent upon its observed value at year k , and thus the mean of Z_{p_L} cannot be calculated for years two through 10. Therefore, Z_{p_L} is excluded as a control variate from years two through 10.

The net supply control variate ($Z_{NS,k}^E$) is used in years $k > 1$ to account for the variability in the number of transfers that occur in the exercise month (t_5). However, the notation for the net supply control variate is changed to $Z_{NS,1}^E$ (where the superscript E and subscript 1 denote the early months (t_0 to t_4) and year one, respectively). The second control variate is used in years two through 10 and adapts the net supply used in $Z_{NS,k}^E$ to control for cost volatility arising from the variability of the number of leases purchased at the beginning of the year (t_0). To account for the variability in the quantity of leases

purchased at t_0 of year k , one examines the net supply that accrues to the city in the latter portion of year $k-1$ (months t_5 to t_{11}),

$$Z_{NS,k}^L = \sum_{i=t_5}^{t_{11}} N_{R,i,k-1} - D_{i,k-1} . \quad [2.38]$$

The CV method, as applied to the multi-year model, is

$$Cost_{var,1} = Cost_1 + c_{1,1} \cdot (Z_{p_L} - E[Z_{p_L}]) + c_{2,1} \cdot (Z_{NS,1}^E - E[Z_{NS,1}^E]), \quad [2.39a]$$

and

$$Cost_{var,k} = Cost_k + c_{2,k} \cdot (Z_{NS,k}^E - E[Z_{NS,k}^E]) + c_{3,k} \cdot (Z_{NS,k}^L - E[Z_{NS,k}^L])$$

for $k = 2, 3, \dots, 10.$ [2.39b]

Control variates are summarized in Table 2.1 including an explanation of how each relates to increases in output variability and the primary factors that influence that variability.

Table 2.1. Summary of control variates used in multi-year scenarios

Control Variate	Applied in Year(s)	Uses...	... to Account for Volatility Arising from...	...Which is Influenced by...
Z_{p_L}	1	The lease price at t_0	Price variability of leases purchased at the beginning of the year (t_0)	Reservoir storage in t_0 .
$Z_{NS,k}^E$	1 – 10	The rate of new water allocation (net supply) from t_0 to t_4 in year k	Variability in leases purchased and/or options exercised in t_5 in year k	Volume of water available to city in t_5 .
$Z_{NS,k}^L$	2 – 10	The net supply from months t_5 to t_{11} in year $k-1$	Variability in quantity of leases purchased in t_0 of year k	Volume of water available to city in t_0 .

2.2.7 Study Region

The U.S. side of the Lower Rio Grande Valley (LRGV) derives its water supply almost entirely from the Rio Grande, with flows managed via the Falcon and Amistad reservoirs (Figure 2.3). The two reservoirs have a combined storage capacity of approximately 5.8 million ac ft (MAF), with an additional 2.1 MAF of capacity set aside for flood protection (dead storage is roughly 30,000 ac ft). The storage in these reservoirs is strictly divided between the United States and Mexico according to the treaty of 1944 [Schoolmaster, 1991], with each country's share of storage, inflows, outflows, and losses calculated as single system-wide values (Table 2.2). Since the two reservoir came on line in 1968, combined U.S. storage in these structures has varied from a low of approximately 0.7 MAF to a high of 4.0 MAF. The hydrologic data record extends from 1970 to 2002, and while there have been subtle shifts in the purpose of the diversions over that period (municipal use increased from 7% to 13% of regional total), average annual usage and monthly usage patterns have remained largely unchanged. The U.S. share of reservoir inflows is allocated to the LRGV's nearly 1600 water rights holders by the Rio Grande Watermaster's Office, which also administers transfers between rights holders.

Ideally, the simulation described would be developed using long time series data sets that cover the same period for each hydrologic parameter (e.g., inflows, outflows), such that serial correlation in and between the data could be preserved. In cases where serial correlation is strong, expected supply and expected demand values would be estimated using conditional probability distributions based on current conditions (or those

in the immediate past). In this case, however, the hydrologic data set is relatively limited (32 years) and use of only the sequential record would have reduced the analysis to a fairly narrow set of conditions. Attempts to expand consideration to a wider range of conditions by fitting existing hydrologic data to standard population models (e.g., lognormal, log-Pearson type III) using chi-square tests yielded very poor fits. The level of serial correlation in data sets and potential relationships between data sets were also explored to determine what other methods of hydrologic input could be used within the simulation.



Figure 2.3. Map of Lower Rio Grande Region

The Pearson test for serial independence was applied to the inflow time series, yielding evidence of weak autocorrelation in the monthly inflow data using both a 1- and 2-month lag (R^2 of 0.15 and 0.05, respectively). The relatively low level of serial dependence is likely a function of the longer time step (i.e., monthly), as well as the arid nature of the watershed and its lack of features that might enhance the system's hydrologic "memory" (e.g., snowpack/snowmelt). Autocorrelation in monthly data is therefore unlikely to play a significant role in simulating regional supply conditions, particularly given that the Valley's regional reservoir capacity is approximately 4 times average annual inflow. This capacity is sufficiently large that ignoring the weak autocorrelation in the data is unlikely to significantly affect simulated reservoir levels, and while interannual correlation of inflows could be an issue in multiyear simulations, it is not a factor in the single annual cycles evaluated in this work. A similar evaluation of the 10-year record of monthly municipal usage (normalized by population) yielded a statistically significant, but weak serial correlation using a 1- and 2-month lag.

Table 2.2. Simulation Data Summary

	Mo.	Jan	Feb	Mar	Apr	May	Jun	Jul	Aug	Sep	Oct	Nov	Dec
Reservoir Inflows (i_t) (x1000 AF)	Mean	89.6	88.5	91.0	100.2	142.2	159.0	159.8	195.8	246.8	203.2	106.3	87.5
	Max	177.8	213.8	156.2	372.4	450.4	410.3	837.1	1095.9	1660.7	748.7	329.5	180.5
	Min	18.3	39.3	44.8	33.0	3.4	51.6	36.3	44.1	58.6	48.7	43.5	36.0
Reservoir Outflow (o_t) (x1000 AF)	Mean	82.3	70.0	94.5	143.2	159.6	152.5	124.7	132.8	97.4	100.6	61.7	59.5
	Max	165.4	175.2	197.8	253.9	345.0	336.5	224.7	739.3	535.0	566.1	157.7	220.3
	Min	13.6	22.4	24.6	6.8	32.7	19.6	19.4	25.7	13.8	0.5	11.0	21.8
Reservoir Losses (l_t) (x1000 AF)	Mean	17.4	21.5	34.3	41.9	46.0	52.6	57.7	55.0	40.8	33.0	22.4	16.9
	Max	25.9	36.4	52.1	59.6	70.6	76.5	86.9	86.8	62.8	51.7	33.3	28.3
	Min	7.7	10.9	19.6	19.5	21.8	24.1	28.1	28.5	20.3	15.0	11.6	9.2
Spot Lease Prices ($p_{L,t}^a$) (\$/ac-ft)^a	Mean	17.0	17.4	16.8	14.6	16.2	16.7	15.2	12.7	15.8	13.8	14.4	16.3
	Max	25.0	30.0	45.0	35.0	30.0	35.0	25.0	20.0	25.0	20.0	20.0	25.0
	Min	10.0	10.0	7.2	10.0	10.0	10.0	10.0	10.0	10.0	10.0	11.0	10.0
Spot Lease Prices ($p_{L,t}^b$) (\$/ac-ft)^a	Mean	27.9	28.5	27.6	26.2	28.0	25.3	23.4	23.5	26.7	25.0	24.9	24.4
	Max	50.0	55.0	50.0	55.0	60.0	50.0	75.0	60.0	55.0	50.0	50.0	55.0
	Min	6.8	7.2	7.2	7.2	7.0	7.0	7.0	7.0	6.8	7.2	6.8	6.8
Demand (d_t) (ac-ft)	Mean	1569	1457	1681	1714	1919	1957	2073	2075	1692	1639	1547	1572
	St.Dev	178.9	195.9	179.7	270	376	383.6	349.6	283.8	299	185.5	193.6	135

^aReservoir Data reflects the years 1970-2002 (IBWC 2004)

^bSpot lease prices reflect the years 1994-2003 (n = 1514)

With respect to relationships between variables, little evidence of correlation was observed between reservoir inflows and municipal water usage ($R^2 = 0.12$, as measured by the Spearman test for trend), a situation that is likely due to climatic differences between central Mexico, where the majority of inflows originate, and the Valley, which is

hundreds of miles away on the Gulf Coast. Correlations were also weak between reservoir outflow and municipal usage ($R^2 = 0.18$), as outflow is dominated by irrigation releases, which in the Valley's semiarid climate are largely dependent on a fixed schedule and tend to obscure the relatively small amount directed to municipal use. These analyses suggest that assuming independence in monthly values for inflow, outflow, and municipal usage could provide a reasonable basis for simulating regional conditions. As a result, values for these variables are randomly selected from the appropriate monthly data list within the simulation. Values for expected supply and expected demand are also computed directly from these monthly distributions (as opposed to conditional probability distributions predicated on current conditions).

Allocations to regional rights holders [2.7] are calculated using an instream loss factor (l_i) of 0.175, and distributed pro rata across the region's 1.9 million ac ft of water rights (N_R). As the number of regional rights substantially outstrips the annual average volume of available reservoir inflows, each acre foot of rights is allocated around 0.725 ac ft of water in an average year. December (initial) reservoir storage levels (R_0) are varied across historical December levels ranging from 0.8 to 2.2 MAF. The city's share of this storage at the beginning of the year (N_{r_0}) is specified as a fraction of the total rights that the city holds (f_{R_0}), such that $N_{r_0} = f_{R_0} \bullet N_{R_T}$. While it might seem logical to assume that high/low levels of R_0 and f_{R_0} would coincide, this is not necessarily the case. A substantial percentage of annual inflows occur in the fall, so even when year-end storage is below average, fall allocations can result in a city beginning the year with a significant volume of carryover water. Three values are chosen to represent low, normal, and high values for both f_{R_0} (0.1, 0.3, 0.5) and R_0 (0.8, 1.5, 2.2 MAF), and paired

combinations of these values represent initial conditions for each simulation. The city's water demand is based on usage records for Brownsville, Texas, a town of 120,000 using an average of approximately 21,000 ac ft per year (Table 2.2).

The vast majority (85%) of regional water use is agricultural, much of it directed toward relatively low valued irrigation activities (e.g., cotton), and a growing municipal population (expected to double by 2050) provides substantial economic incentives for agricultural to urban water transfers. While economic incentives alone do not always translate to an increased volume of trades (DeMouche, Ward et al. 2003), this does appear to be the primary driver in the Valley (Chang and Griffin 1992). The regional water market is relatively efficient and has presided over the steady transfer of permanent rights from irrigators and urban users in recent years (Griffin 1998). Permanent transfers are almost always approved but must navigate a regulatory process that can take over a year to complete. Leases tend to raise fewer concerns over third-party impacts and are subject to a simplified approval process that is often concluded in a few days (Griffin and Characklis 2002). Lease transactions require only that the buyer and seller deliver a one-page document to the watermaster detailing their respective account numbers and the volume of water to be transferred (price information is optional). The ease of completing these transactions contributes to the high level of market activity, with an average of nearly 70,000 ac ft of water transferred via leases each year (Watermaster's Office 2004). The structure of the market leads to the assumption that spot market transaction costs are essentially negligible. While this assumption is reasonable within the Valley, it may not be so in many other regions, a factor which may bias this analysis in favor of spot market leases.

All water markets exhibit idiosyncrasies. In the case of the LRGV, the most noteworthy is that current rules allow for permanent rights to be transferred between agricultural and urban users, but only allow lease transactions between similar user types (e.g., urban to urban), giving rise to two spot lease markets [see Characklis et al., 1999]. The municipal market has fewer transactions, as cities tend to hold volumes of permanent rights well in excess of average usage, while the agricultural lease market is much more active (1514 transactions over the period 1994–2003; average price \$22.60 per ac ft). Efforts to eliminate this prohibition on intersectoral leasing are currently being undertaken (Group 2000), and when this occurs it seems likely that the lower marginal value of irrigation water will lead to regional lease prices similar to those observed in the agricultural market. These simulations assume this is the case and that lease prices from the agricultural market are representative of what would be observed in agricultural to urban transactions.

An analysis was undertaken to explore statistical correlations between spot lease prices and several hydrologic parameters (e.g., reservoir storage, inflows, outflows), the idea being that if a low reservoir level in December (when options are bought) is a strong indicator that spot market prices in May (when options are exercised) will be higher, a well-informed market would incorporate consideration of this into option/exercise prices. Results suggest that the only parameter exhibiting significant explanatory power over lease prices is reservoir storage, but linear correlations between lease price and storage levels yield very weak predictive relationships. Further analysis using the Wilcoxon two-sample test strongly indicates ($p\text{-value} < 0.0001$) that there are two separate populations of lease price data, one when reservoir storage is above 1.43 MAF and another when

storage is below this level. Monthly lease price data are therefore separated into two lists based on observations made when reservoir levels are either above ($\hat{p}_{L_t}^a$) or below ($\hat{p}_{L_t}^b$) this threshold (Table 2.2).

It should also be noted that while there is some evidence of serial correlation (again using the Pearson test) in the spot price data set as a whole, once the data are separated into these two subsets the effects of serial correlation becomes quite weak (1-month lag typically has an $R^2 < 0.10$). In effect, it appears that when reservoir storage drops below (rises above) the threshold level, the mean monthly price increases (decreases), but subsequent price variation about the mean is essentially random. This randomness in spot market prices is likely due, in large part, to the decentralized nature of the market. While the prices of the most recent lease transactions can be obtained from the watermaster's office, it seems clear that most transactions are completed with only a general knowledge of the current level of water scarcity (i.e., reservoir level is low or its not). This leads to a spread in prices, even those observed in the same month with similar reservoir levels. Such behavior might suggest that a high-volume buyer, motivated by large potential savings, could find a lower price by increasing the amount of time and effort spent looking for a seller. However, correlations between spot market prices and the volume purchased yielded no evidence of a statistically significant relationship. Finally, consideration was also given to adjusting the spot price data to reflect real prices over the period 1993–2002. Both the producer price index for all farm products (which rises from 106.3 to 111.5 over this period) and the Texas index of prices received for farm products (which falls from 98.0 to 93.0 over the same period) seem likely to be strong indicators of variation in the marginal benefits of irrigation water over

time, but the mixed directions and small changes in these indices led to the decision to use unadjusted (nominal) lease prices. Chi-square tests yielded little evidence that monthly lease data fit any standard distribution type, so lease prices are represented as monthly data lists. The simulation is set up to randomly sample spot prices from one of monthly lists, with the decision as to which made according to the current storage level are sampled monthly from the appropriate list.

Option contracts have been discussed in the LRGV but are not yet actively traded. Their introduction into the market, however, would appear to be a logical step with few bureaucratic hurdles. Within the simulation, a single type of European call contract is offered, with the option purchased on 31 December ($t = 0$) and exercised on 31 May ($t = 5$). The date 31 May falls just before the peak usage months in both the municipal and agricultural sectors and therefore seems to provide a logical point for users to assess their current supplies and make choices. There are, of course, a host of other call dates that might be suitable as well, and consideration might even be given to developing option contracts with multiple exercise dates, but such considerations are left for future work. Given an initial reservoir storage (R_0), the conditional probability of May storage (R_5) being above or below 1.43 MAF can be computed, and it is assumed that the market incorporates this information into option pricing. As a result, equation [2.1] is modified such that the option price is conditional on R_0 , with

$$\begin{aligned}
 P_{O|R_0} = & e^{-r \cdot 5} \cdot P[R_5 \geq 1.43\text{MAF} | R_0] \cdot E \left[\max(\hat{p}_{L_5}^a - p_x, 0) \right] + \\
 & e^{-r \cdot 5} \cdot P[R_5 < 1.43\text{MAF} | R_0] \cdot E \left[\max(\hat{p}_{L_5}^b - p_x, 0) \right].
 \end{aligned}
 \tag{2.40}$$

The exercise price (p_X) is set at \$15 per ac ft, a level in line with the mean spot lease price when reservoir level is above the threshold level and one that is therefore assumed to be sufficient to attract enough irrigators to create an options market. Using this exercise price, the resulting option prices are \$13.26, \$11.36, and \$2.18 per ac ft when initial storage levels are 0.8, 1.5, and 2.2 MAF, respectively. The annualized price of permanent rights (p_R) is \$22.60 per ac ft, but considering that only about 0.7 ac ft are allocated to these rights in an average year, the effective annualized cost of water obtained via these rights is \$31.17 per ac ft. The annualized cost of rights corresponds to a \$1000 per ac ft purchase price amortized over 40 years at a 6% discount rate, and assumes that the real value of the right increases at around 4% per year over that period.

2.3 RESULTS

2.3.1 Developing Minimum Cost Portfolios

All portfolios are developed with respect to a 1-year planning horizon using the least favorable set of initial conditions ($f_{R_0} = 0.1$; $R_0 = 0.8$ MAF), with minimum cost portfolios identified for strategies A (permanent rights alone), B (rights and options), and C (rights, options and leases) (Figure 2.4). Several reliability levels are assessed, with reliability defined relative to the initial conditions. In other words, a portfolio providing 99.5% reliability under the least favorable conditions would translate to an even higher reliability if the same portfolio were used under better conditions. Critical failures are limited to $<0.5\%$ in all cases.

Achieving 99.5% reliability using permanent rights alone (A) requires the maintenance of just over 70,000 ac ft of rights with an annual cost of \$1.59 million. The volume of permanent rights is fixed throughout the year, so this cost is invariant, but reducing reliability from 99.5 to 99% lowers expected costs by \$0.1 million (Table 2.3). Reducing reliability from 99 to 98% lowers annual costs by \$0.09 million, indicating that the marginal cost of reliability rises with increasing reliability. Most failures occur in December, but on average there is a substantial volume of water leftover at year's end (23,200 ac ft).

Using strategy B, a 99.5% reliability level can be achieved with 53,000 ac ft of permanent rights and 11,000 ac ft of options (4900 ac ft of which are exercised on average). The expected annual cost of this portfolio is \$1.34 million, a reduction of a little over \$0.25 million (16%) relative to strategy A. The ability to make acquisition decisions in May, when improved information is available, also leads to a significant

reduction in the average volume of water remaining in the city's possession at year end (17,100 ac ft). This not only reduces the city's expected costs, but also makes more water available to other regional users in most years. Strategy B results in some cost variability, but the interquartile cost range (i.e., the 25th to 75th percentile) extends from only \$1.32 to \$1.35 million. The CVAR is \$1.37 million, small relative to the expected value, indicating that the use of options can significantly reduce expected costs while still limiting the city's exposure to large cost fluctuations. The marginal cost of reliability (\$0.1 million/percentage point from 99% to 99.5%) is approximately half of that for strategy A, but the marginal cost increases for both strategies as reliability rises.

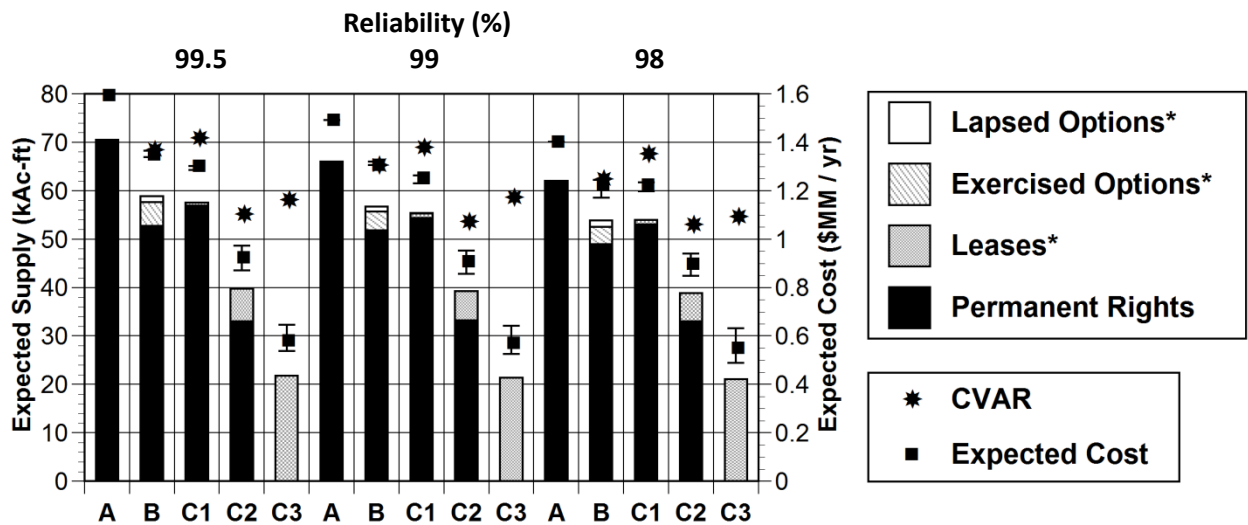


Figure 2.4. Minimum cost portfolios for different strategies ($f_{R_0} = 0.1$; $R_0 = 0.8$ MAF). Dollars are in millions.

The volume of permanent rights in strategy B is driven largely by the monthly allocations required to reliably meet demand prior to 31 May when options can be exercised. In this case, if permanent rights were reduced below 53,000 ac ft, the number of failures occurring before the city could exercise would make it impossible to maintain an overall reliability of 99.5%. With only rights and options, the city has one opportunity

to augment its supply during the year; consequently, the values for α_2 (1.67) and β_2 (1.85) must be relatively high to ensure that the 99.5% reliability goal is met (Table 2.3). The value of α_2 declines with lower reliability as the city allows the ratio of expected supply to expected demand to drop a little lower before acquiring

Table 2.3. Minimum Cost Portfolios^a

Reliability	Strategy	Expected Cost, millions of dollars	CVAR, millions of dollars	α_1	β_1	α_2	β_2	Expected Year-End Supply, ac ft x 1000
99.5%	A	1.59						23.2
	B	1.34	1.37			1.67	1.85	17.1
	C1	1.30	1.41			1.30	1.31	15.6
	C2	0.92	1.10	1.56	1.77	0.93	1.04	7.1
	C3	0.58	1.16		2.56	0.97	1.09	2.4
99%	A	1.49						20.4
	B	1.30	1.31			1.48	2.10	14.3
	C1	1.25	1.38			1.20	1.28	14.1
	C2	0.91	1.07	1.48	1.74	0.90	0.93	6.5
	C3	0.57	1.16		2.50	0.96	1.04	2.1
98%	A	1.40						17.9
	B	1.23	1.25			1.33	2.15	13.5
	C1	1.22	1.35			1.20	1.23	13.4
	C2	0.90	1.06	1.62	1.69	0.70	0.79	6.1
	C3	0.55	1.09		2.32	0.75	1.07	1.8

^aAll portfolios assume an initial reservoir storage (R_0) of 0.8 MAF and an $f_{R_0} = 0.1$.

more water. Meanwhile, β_2 rises from 1.85 to 2.15 as reliability declines, suggesting that when the city does exercise options, it will exercise slightly more. It should be noted, however, that in this case the expected costs are not very sensitive to small differences in the β_2 . Once β_2 is sufficiently large to ensure that enough options are exercised to meet reliability goals, then small increases in its value only lead to a few more options being exercised and an almost imperceptible increase in expected costs. For example, in the case of 99% reliability, varying β_2 from 2.10 to 2.40 increases the volume of exercised

options from 3888 to 3956 ac ft and raises expected costs less than a thousand dollars. By contrast, a similar variation in α_2 would have a greater impact on expected costs as it would increase the number of acquisitions made, not just their size. Expected costs would also be more sensitive to variation in β_2 if the number of options the city holds were higher. In some situations the solution surface is quite flat in the neighborhood of the expected cost minimum, and randomness in the search path can lead to the identification of portfolios with nearly identical values for expected cost and reliability but different α and β values. The guidelines set for the simulation and search algorithm provide a resolution that was deemed appropriate for this work, but this resolution could be further sharpened at the cost of increased computation time.

Strategy C involves consideration of permanent rights and both types of temporary transfer (options and leases). In this case, opportunities for spot market acquisitions, in combination with the relatively high costs of permanent rights, would lead a city interested solely in minimizing expected costs to eliminate permanent rights from its portfolio. Such a strategy provides an interesting lower bound but is unlikely to be widely adopted, so several alternative portfolios are considered:

In portfolio C1, the city is willing to use temporary transfers to reduce its supply costs but is concerned over the risks associated with cost variability and will not accept a portfolio for which the CVAR exceeds expected costs by more than 10% (i.e., constraint [2.23] is employed).

In portfolio C2, the city maintains 33,000 ac ft of permanent rights, an amount that will yield a little more water than the city's average annual demand of 21,000 ac ft in most years (although timing between supply and demand may not coincide). The city

then considers the use of temporary transfers to supplement its supply but places no limits on cost variability.

In portfolio C3, the city maintains no permanent rights and relies entirely on temporary transfers to meet demand and places no limits on cost variability.

Limits on the CVAR to expected cost ratio result in the C1 portfolio depending primarily on permanent rights (56,900 ac ft) with a small volume of spot market leases (but no options) used to augment supply. Expected costs decline only slightly relative to strategy B, while the CVAR rises but remains within the imposed limit. There is also a small decline in the average year-end supply. The large volume of rights ensures the city will not need to resort to the spot market before May, so α_1 and β_1 are not applicable. In the latter portion of the year the α_2 value indicates that the increased acquisition opportunities allow the city to be less risk averse than with strategy B, waiting until the expected supply to expected demand ratio drops to 1.30 instead of 1.67 before acquiring water (β_2 drops to 1.31, indicating that acquisitions are also smaller). Decreasing α_2 and β_2 serves to lower reliability, with the marginal cost of reliability remaining relatively similar to that of strategy B from 99.5 to 99%.

Expected cost drops significantly using strategy C2 (\$0.92 million at 99.5% reliability). This is accompanied by a CVAR of \$1.10 million, which is substantially less than that observed for strategies A, B, or C1 but still pushes the CVAR to expected cost ratio up to 1.20. There is also a considerable decrease in the average volume of water leftover at year's end (7100 ac ft). Options again play no role, as the greater flexibility of the spot market and lack of concern over the CVAR make leasing a less expensive means of meeting reliability constraints. The increased flexibility of the spot market also results

in lower marginal costs for reliability. The unfavorable initial conditions result in the expected supply to expected demand ratio being quite low at the beginning of the year. Therefore, as long as α_1 is set above this level, small variations in its value will have little impact on reliability (i.e., the city will always buy at the beginning of the year unless α_1 were set very low). Small changes in acquisition size (β_1), however, will lower reliability. In this case, the relatively large acquisitions made in December provide enough water so that post-April acquisitions are smaller (i.e., $\beta_2 < 1$) and made when the supply to demand ratio is quite low (i.e., $\alpha_2 < 1$); thus they serve as a means of subtly adjusting supply in the latter part of the year.

The expected cost of meeting 99.5% reliability through strategy C3 declines to \$0.58 million with a portfolio that relies entirely on spot market leases. Dependence on spot leases results in a CVAR that is roughly twice the expected cost, although still lower than the expected cost of the A, B, and C1 strategies. The city begins the year with no permanent rights and will need to buy water immediately, so α_1 values are meaningless. The high β_1 (2.56) points to a large acquisition in $t = 0$, large enough that only subtle adjustments to supply are required over the remainder of the year to meet reliability objectives in most years. In this case, the size of the initial acquisition and the fact that it is always the same size (for a given set of initial conditions) mean that most of the variability in portfolio cost is due to price volatility, not differences in the timing or magnitude of acquisitions. This leads to an interquartile range that is narrower than might be expected. The range is still considerably wider than that of C2 in relative terms, since C2's expected costs are 60% higher, but in C3 a much larger fraction of annual demand is met with this initial acquisition. Very dry years still result in large late-year

acquisitions and lease prices in December can be high in some years, both of which contribute to the large CVAR, but in at least half the years, annual costs will fall within $\pm 12\%$ of the expected value.

When considering the practicality of each strategy, the realities associated with managing a utility make it unlikely that strategy C3 would be widely adopted. Furthermore, the increase in CVAR that occurs when switching from C1 to C2 (\$0.31 million at 99.5% reliability), would seem a small price to pay for the significant reduction in expected costs. This leaves strategies B and C2 as perhaps the most attractive alternatives to sole reliance on permanent rights, given that both significantly reduce expected costs while limiting a city's exposure to wide cost swings. As a result, these two strategies receive further analysis under a broader range of initial conditions. Table 2.4 describes minimum cost portfolios (99.5% reliable) developed using strategies B and C2 under more favorable initial conditions. Portfolios are most sensitive to changes in the initial water supply (f_{R_0}), with the expected cost of strategy B ($R_0 = 1.5$ MAF) declining from \$1.32 to \$0.66 million as f_{R_0} rises from 0.1 to 0.5, respectively. The portfolio developed using strategy C2 maintains greater flexibility through the use of spot leases, so while costs decline with rising f_{R_0} , the change is relatively small (of the order of \$0.01 million). Changes in initial reservoir storage (R_0) affect only the price of options and leases (not the amount that must be bought), and while higher initial storage levels result in slightly lower expected costs, there is little impact on portfolio composition. The expected costs of a portfolio using strategy B ($f_{R_0} = 0.1$ or 0.3) decline approximately \$0.1 million as R_0 rises from 1.5 to 2.2 MAF and the effects on a portfolio developed using strategy C2 are even smaller.

Table 2.4. Strategy B and C2 Portfolios Under Different Initial Conditions (R_0, f_{R_0})^a

Initial Reservoir Level, R_0 (MAP)		Strategy B						Strategy C2					
		f_{R_0}	N_R , ac ft x 1000	N_0 , ac ft x 1000	N_X^b , ac ft x1000	Cost, ^b millions of dollars	CVAR, millions of dollars	Year End Supply, ^b ac ft x1000	N_R , ac ft x 1000	N_0 , ac ft x 1000	N_L^b , ac ft x 1000	Cost, ^b millions of dollars	CVAR, millions of dollars
1.5	0.1	52.9	6.1	3.4	1.32	1.36	16.0	33.0	...	7.1	0.87	0.98	7.5
	0.3	34.2	2.5	1.6	0.83	0.84	7.5	33.0	...	0.6	0.76	0.84	7.7
	0.5	29.0	0.0	0.0	0.66	...	6.2	33.0	...	0.0	0.75	...	13.8
2.2	0.1	53.0	6.1	3.3	1.26	1.30	15.8	33.0	...	6.7	0.85	0.94	7.3
	0.3	27.1	7.7	6.5	0.70	0.72	7.5	33.0	...	0.7	0.76	0.81	7.7
	0.5	29.0	0.0	0.0	0.66	...	6.2	33.0	...	0.0	0.87	...	13.8

^aAll portfolios 99.5% reliable.

h—

Both strategies B and C2, regardless of initial conditions, are expected to leave the city with at least 30% of its average annual water supply (21,000 ac ft) available for use in the next year. The same applies for any of the strategies described in Table 2.3, with the exception of C3 (all water obtained via spot market), implying that even though this analysis is limited to a 1-year horizon, the approach is not likely to generate portfolios that will leave the city in an untenable position at year's end (i.e., without any water). The approach described in this paper may therefore provide a reasonable starting point for future work seeking to develop long-term portfolios.




Figure 2.5. Variations on strategy C2.

Varying the relative distribution of leases and options provides a means of “fine tuning” the trade-offs between expected costs and cost variability. Besides limiting cost variability, options can also provide some practical advantages in multi-year planning as they provide an opportunity for long-term revolving contracts. These might involve the city making an annual payment for a specified volume of options each year. Such a

contract could limit the city's exposure to spot market volatility, while still allowing some access to the flexibility the spot market provides. Figure 2.5 describes a range of variations on strategy C2, each containing 33,000 ac ft of rights and meeting 99.5% reliability through various combinations of leases and purchased options. Under the least favorable initial conditions, a city could reduce its expected number of leases 25% (from 6860 to 5270 ac ft) with a contract for 4000 options, resulting in a portfolio with expected costs only slightly higher (\$0.05 million) than one without options, but with a somewhat lower CVAR (\$1.12 versus \$1.095 million). While the reduction in CVAR is modest, it should be noted that there are additional benefits that might be associated with some form of long-term option that are not included in this analysis. When either transaction costs or transaction risk are relevant factors, long-term option contracts are likely to become increasingly attractive relative to spot market leases, but quantifying these values is difficult. As part of a long-term contract, the city would be committed to the option payment during years in which conditions were more favorable, but it would be less vulnerable to large swings in lease price during other years. While an assessment of multiyear strategies is beyond the scope of this work, it does appear that some variation of C2 might serve as a foundation for a city seeking to lower long-term water supply costs through the use of multiyear option contracts. Annual increases in the number of permanent rights could be made to keep base capacity in line with demand growth, while long-term option contracts could reduce the need for leasing while providing added security and insulation from large swings in spot market prices.

2.3.2 Application of CV Method to Single-Year Model

Figure 2.6a plots the expected annual portfolio cost ($Cost$), while Figure 2.6b plots the variance-reduced cost ($Cost_{var}$). Each expected value plotted in both figures was created using 125 simulation runs (i.e. realizations); however, in the variance reduced case (Figure 2.6b) 25 of these are dedicated to the pilot study in order to calculate c^* values. These values are then applied to the 100 realizations performed in the main model (Figure 2.2) to calculate the variance reduced estimate of expected cost ($Cost_{var}$).

The smoothing effect of the CV method improves the accuracy and efficiency of the optimization of the model through noise reduction. Noise can be measured as standard error [2.24], but here, the definition is slightly modified to be the average standard error ($\overline{s.e.}$) as a percentage of the average portfolio cost (\overline{Cost}),

$$Noise = \overline{s.e.} / \overline{Cost} \cdot 100. \quad [2.41]$$

Figure 2.7 provides a comparison between the average optimal values of portfolio cost produced with and without the CV method using the same three levels of computational effort (125, 150, and 300 total realizations per “call” from the implicit filtering search algorithm, or “optimizer”) and ordered according to the average noise level. For the runs utilizing the CV method, the pilot study accounted for 25, 25, and 50 of the total number of realizations, respectively. In stochastic optimization, search algorithms rarely converge to a unique solution, but rather to a relatively circumscribed region. As such, in Figure 2.7 each data point reflects the average optimized value of 100 optimization runs with the error bars representing the 25th and 75th percentiles, a range that varies from 1% to 4% of the mean, depending on the number of realizations.


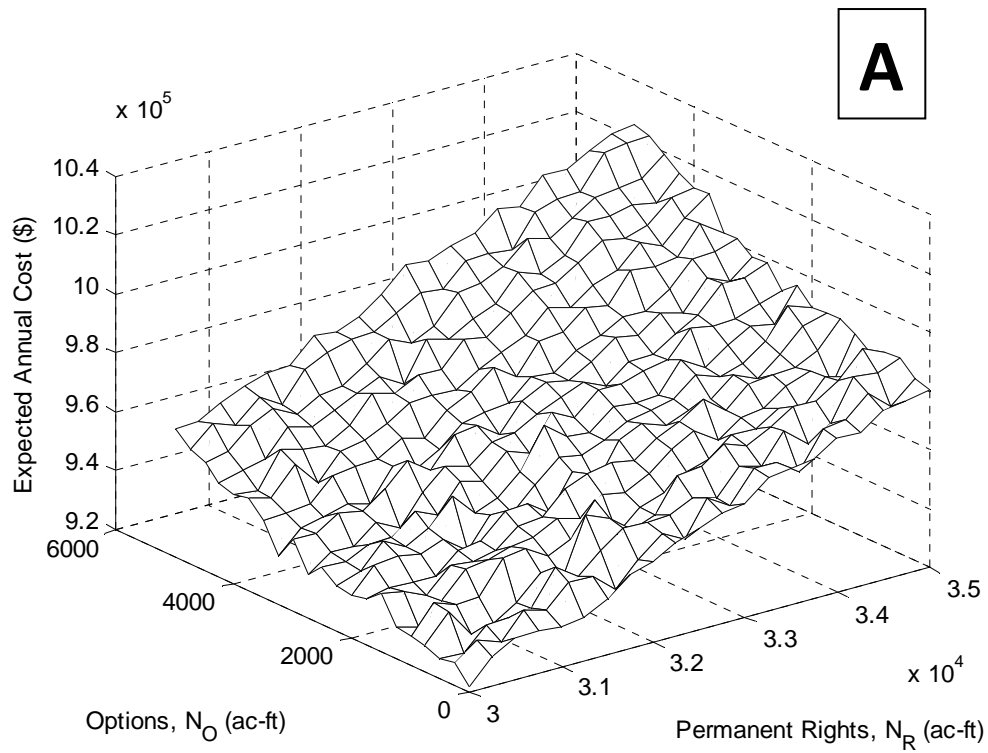


Figure 2.6. Portfolio landscapes produced using 125 total realizations, with (B) and without (A) the CV method.

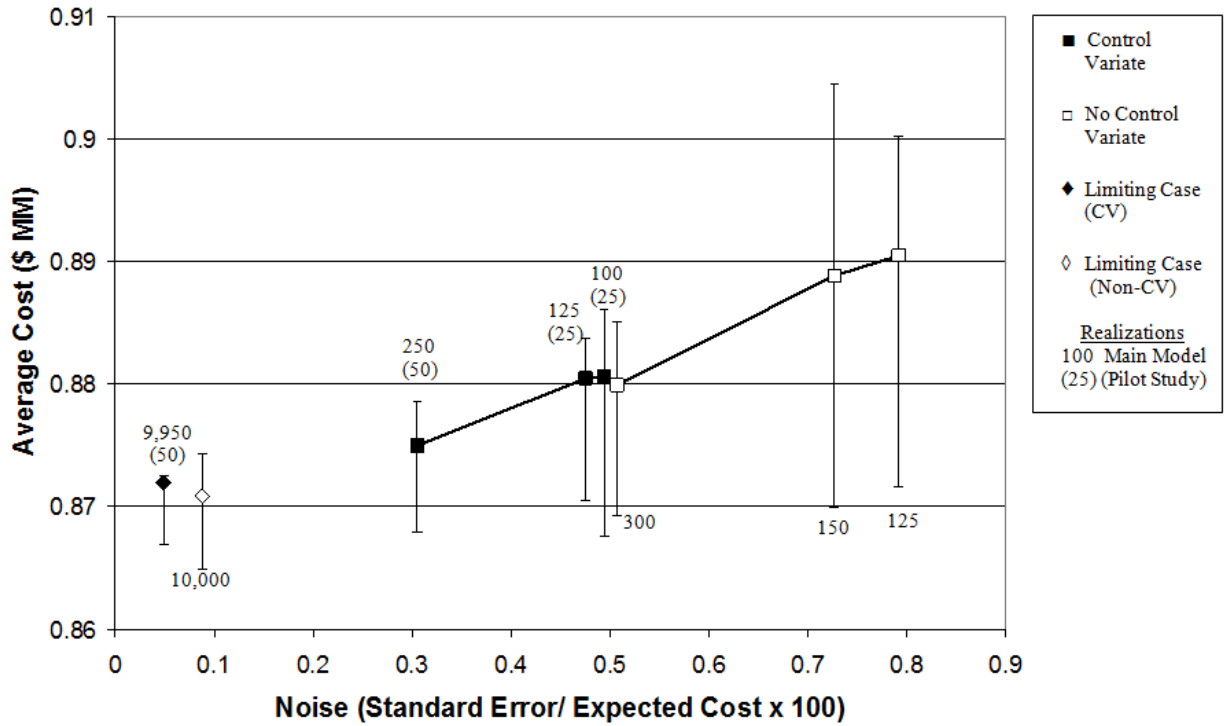


Figure 2.7. Average optimized results of the one-year model produced when controlling for noise, with and without the CV method. The scenario allowed the use of permanent rights, options, and leases, and initial conditions set to $R_0 = 800,000$ ac-ft and $f_{R_0} = 0.1$. Note: Error bars represent 25th and 75th percentiles.

The computational effort required to achieve optimality conditions can be measured either through the amount of work invested in the simulation (i.e. number of realizations) or through the number of times the optimizer calls the simulation. In this case the number of calls from the optimizer is not significantly affected when the CV method is applied. Instead, the application of the CV method allows for the same level of accuracy and precision in the objective function value to be achieved with a significant reduction in the number of realizations required each time the optimizer calls the simulation. Results obtained both with and without the CV method, show that the more times the simulation is run (i.e. more realizations) the lower the average optimized value.

Also, the range of optimal values arising from multiple optimization attempts (as reflected in the error bars) shrinks noticeably as the noise is reduced, evidence that use of the CV method can also improve objective function values in terms of both lower portfolio costs and greater precision. For comparative purposes, two limiting test cases (one with and one without the CV method) estimate a lower bound on the gains possible through noise reduction by using 10,000 realizations per optimizer call, providing an estimate of a lower bound on the improvement in optimal values possible if computational time were not a concern.

The important point in Figure 2.7 is that the CV method allows for essentially the same minimum average portfolio cost (\$0.88MM) to be identified with 125 realizations as would require 300 realizations if the CV method were not used (and with roughly the same level of precision). Over the course of an optimization, one in which the search algorithm were to call the simulator 100 times, the CV method would reduce the total number of times the simulation needed to be run from 30,000 to 12,500, leading to a substantial reduction in computing time.

An examination of one-year portfolios is conducted using the CV method, with all portfolios meeting a minimum monthly reliability of 99% (i.e. one failure every 100 months). Three different strategies are considered (Figure 2.8):

- (A) Permanent rights only. This represents a “typical” case in which the city maintains a sufficient volume of rights (i.e. firm capacity) to meet its reliability objectives.
- (B) Rights and options. Demand is met through a combination of permanent rights and options. The city purchases a constant volume of options each year,

but varies the volume exercised based on supply conditions at the end of the exercise month (May).

- (C) Rights, options, and leases. Demand is met using permanent rights, options and leases. Leases can be acquired at any monthly interval throughout the year.
- (D) Same as strategy C, but ratio of CVAR to expected cost is no greater than 1.1.

The assumption here is that, in cases (B and (C)), the city is unlikely to adopt a portfolio in which firm capacity cannot at least meet demand in a typical year. So, all portfolios maintain a minimum of 30,000 ac-ft of permanent rights, using leases and options as supplements during dry periods. Nonetheless, the use of options and leases can significantly reduce the amount of firm capacity that must be maintained to meet reliability goals, thereby lowering expected costs.

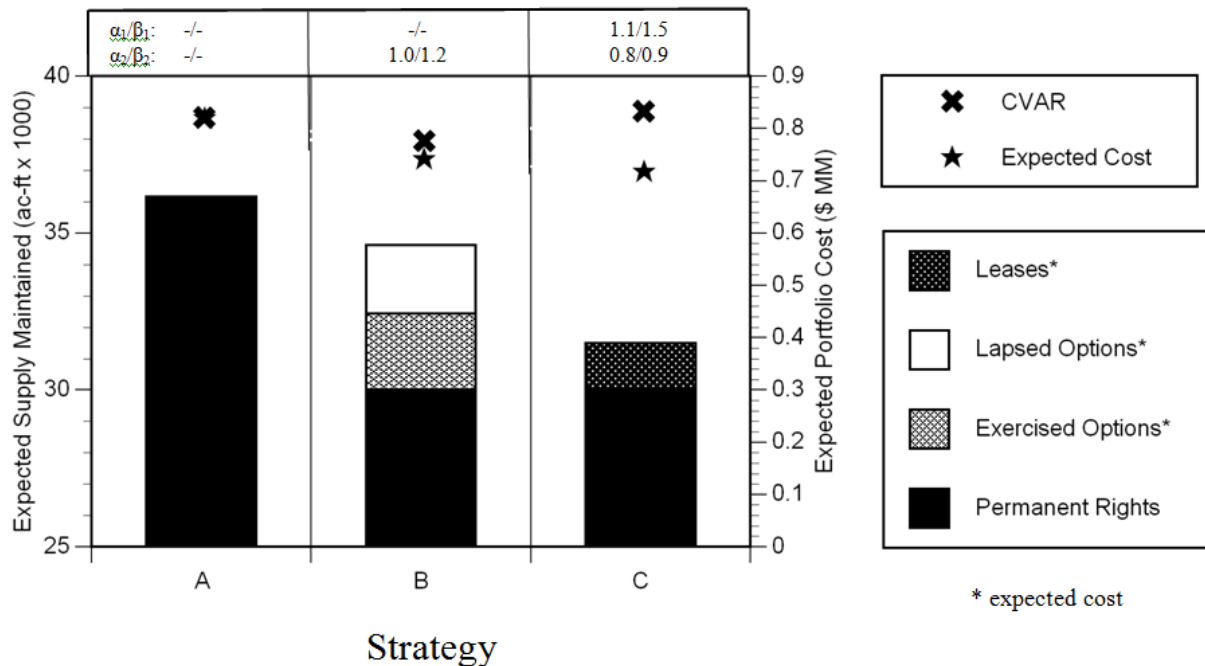


Figure 2.8. Optimized one-year portfolio results ($f_{R_0} = 0.3$).

2.3.3 Multi-Year Portfolio Scenarios

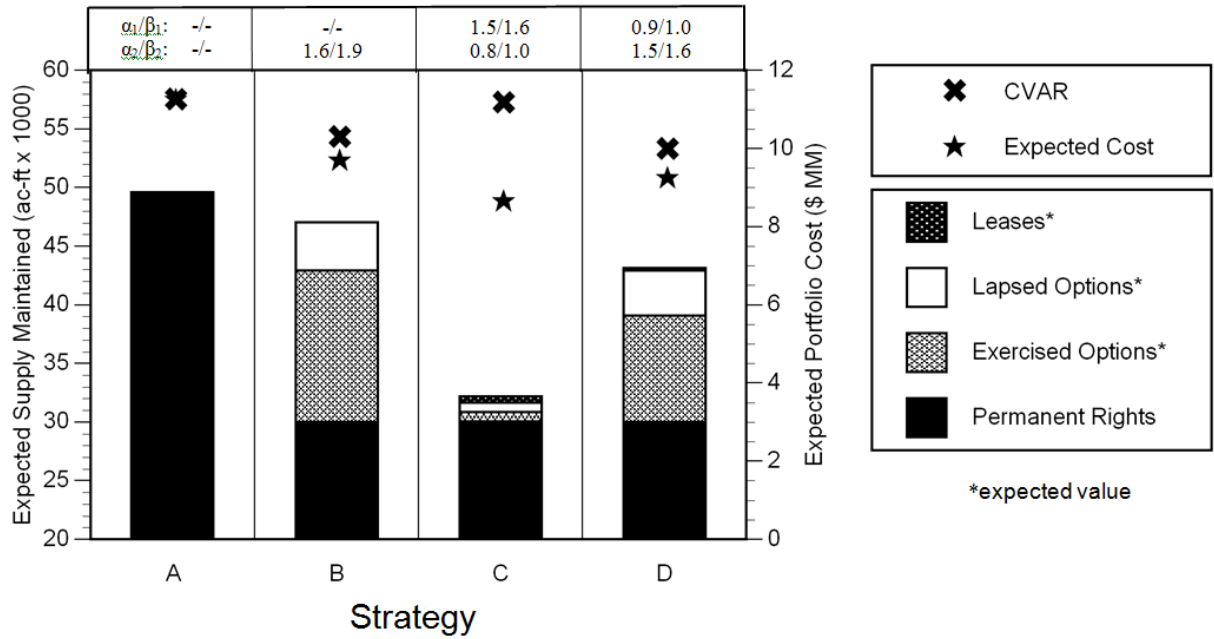


Figure 2.9. Optimized 10-year portfolio results ($f_{R_0} = 0.3$).

The same three strategies, plus one additional strategy, are used to develop portfolios over a 10-year period, with each again meeting a minimum of 99% monthly reliability (Figure 2.9). Initial conditions remain unchanged, but a 2% annual rate of demand growth is assumed over the 10 years. The other important distinction is that the number of options purchased each year is constant, reflecting a long-term contract with option payments due at the beginning of each year, and exercise payments due when/if optioned water is called in May. Strategy A reflects the situation in which firm capacity is maintained, and the city requires nearly 50,000 ac-ft of permanent rights to meet its reliability goals at a total cost of \$11.2 million over 10 years. Allowing options to be purchased in conjunction with owning permanent rights (strategy B), the city is able to

reduce firm capacity and lower its expected 10-year portfolio cost by 14%, down to \$9.7 million. Further savings are realized when leases are also considered (strategy C), lowering the expected portfolio cost to \$8.6 million, but there is greater variability in portfolio cost, as described by the CVAR values. A closer inspection of the data shows that the CVAR-to-expected cost ratio rises from 1.06 in strategy B to 1.29 in strategy C, which highlighting the role that options can play in reducing cost variability by limiting exposure to lease price volatility. Strategy D employs rights, options, and leases, just as in C, but with a cost variability constraint which limits the CVAR-to-expected cost ratio to less than 1.1. While strategy D leads to less cost variability, it costs \$600,000 more over the 10-year period.

The efficacy of the CV method is tested on the ten-year model just as with the one-year model. The average optimized multi-year portfolio costs were evaluated both with and without the CV method using different numbers of realizations (Figure 2.10). Similar to the one-year case, the ten-year optimization results demonstrate that the use of the CV method reduces the number of realizations required to reach a given solution (with an equivalent level of precision) by at least 50%. However, in this case, the savings in computational time is magnified by the longer simulation period, so in a scenario in which the optimizer calls the simulator 100 times, equivalent optimal values could be achieved with 125 realizations per call using the CV method as compared with 300 without (a total savings of 185,000 simulation runs). Perhaps more to the point, using a high end PC it takes roughly 12 hours to optimize a specific portfolio using 125 realizations, whereas optimizing the same portfolio using 300 realizations requires ~30 hours.




Figure 2.10. Average optimized results of the 10-year model obtained controlling for noise, with and without the CV method. The scenario allowed the use of permanent rights, options, and leases, and initial conditions set to $R_0 = 800,000$ ac-ft and $f_{R_0} = 0.1$. Note: Error bars represent 25th and 75th percentiles.

Expected costs over the 10-year period exhibit less noise than those in the one-year case, as a sum of ten years' costs tends to moderate the relative noise arising from extreme events observed in a single year. While this reduces the relative noise reduction possible through use of the CV method, its use still results in a significant improvement in computational efficiency. Nonetheless, this reduction in noise lowering potential may be more of a factor when considering much longer time horizons. Another issue to be considered by future researchers is that the effectiveness of the CV method is not constant across the solution surface, which becomes apparent when expected portfolio costs are mapped across a range of permanent rights and options (holding α/β values constant) (Figure 2.11). The application of the CV method over the domain in Figure

2.11 yields variance reduction ranging from 0% to over 70%. When surveying this solution surface, it is clear that application of the CV method significantly reduces variance across broad sections of the landscape, but two areas of low variance reduction are observed. The first is when the city maintains a large volume of permanent rights (> 45,000 ac-ft), well beyond its average annual demand (21,000 ac-ft), such that the city can reliably meet demand with very few leases or options. A low volume of these transfers corresponds to little variability in expected portfolio costs, therefore the CV method has a limited ability to reduce variance. Variance reduction is also low when the volume of permanent rights and options is very low, a situation in which the city must rely largely on spot leases. While it seems paradoxical that regions involving both more and fewer transfers should both translate to low variance reduction, this region of heavy spot lease activity exhibits some of the greatest levels of noise, which largely arise from variability in the lease price. However, in the 10-year model the lease price variability is controlled for by the CV method only at t_0 of the first year. The other control variates ($Z_{NS,k}^E$ and $Z_{NS,k}^L$) only account for variability in the quantity of transfers executed throughout the simulation period. While regions of low variance reduction do exist, the minimum cost portfolios described in this work are largely located in the broad swath of greatest variance reduction, and are also those likely to require the greatest level of analysis to identify. Portfolios composed mostly of permanent rights require little analysis, while those with a large reliance on spot leases are unlikely to be practical in most cases, variance reduction would appear to be useful in identifying the types of portfolios that will be of greatest interest to utilities.




Figure 2.11. Map of variance reduction achieved across a range of permanent rights and options for 10-year model. ($\alpha_1=1.3$; $\beta_1=1.5$; $\alpha_2=1.0$; $\beta_2=1.15$)

2.3.4 Long-Term Option Contracts

Thus far, results have revolved around static 10-year portfolios, in which the number of permanent rights and options purchased remain constant throughout. While this is the sort of stability that utilities have traditionally sought, given an annual 2% growth in demand, it may be possible to develop multi-year arrangements that accommodate this growth more effectively. Figure 2.12 details the portfolio compositions and costs of three scenarios involving more creative long-term option contracts, with all scenarios based on strategy B (permanent rights and options). The first panel (B) describes the year-by-year evolution of portfolio B described earlier in Figure 2.9, where the city has a long-term contract for a constant number of options. As one would expect, the average number of exercised options increases as demand increases. Alternatively, the second panel (B2) reflects a multi-year option contract in which the

contract involves scaling the number of options purchased each year to correspond with demand growth according to [2.36]. This type of contract involves fewer option purchases in the earlier years, and reduces the expected cost of the option contract by 13%, or \$378,000 over its life (a 5% change in the total portfolio cost). The third panel (B3) structures the option purchase so that at the beginning of each year, the city either purchases a base volume of options, or a larger volume, 30% greater than the base (roughly the same increase observed between years one and 10 of B2). The volume of options purchased is based on the city's water supply entering the new year as measured by f_{R_0} . The base volume of options is purchased if $f_{R_0} > 0.2$, or if the city's current supply is more than 20% of the total volume of rights it holds, otherwise the city purchases the larger volume. This contract (B3) results in a reduction of the contract's expected cost by 23% relative to the static case (B), and a total savings of \$658,000 over 10 years.




Figure 2.12. Expected option activity over 10 years, given (B) static option purchase, (B2) option purchases that grow with expected demand growth, and (B3) two levels of option purchases, dependent upon water supply conditions.

The contingent clauses described in option strategies B2 and B3 should not be difficult to write into a long-term contract, and both provide the city with long-term security and greater flexibility. The option contract in B2 is structured with annual demand growth in mind, however, if that growth does not materialize, B3 might be more advantageous in that the utility is not committed to purchasing the larger volume of options, providing additional flexibility.

2.4 CONCLUSIONS

Many cities with access to water markets currently rely on permanent rights alone to meet demand. The results of this work suggest that expanding a city's water supply portfolio to include options and/or leases could significantly lower expected costs while maintaining high levels of reliability. Considerable reductions in expected cost can be realized through the introduction of options alone, but the use of spot market leases can cut costs even further. While it is unlikely many cities would undertake a supply strategy that relied entirely (or even primarily) on temporary transfers, more conservative approaches in which leases and options supplement a substantial base capacity of permanent rights can still reduce expected costs significantly. While options play a relatively limited role in the portfolios developed in this analysis, some of this is attributable to assumptions regarding the spot market (i.e., no transaction costs, unlimited availability) and a risk-neutral utility. A more risk-averse utility facing a less "liquid" spot market might find options more attractive, particularly when developing multiyear water supply strategies. In addition, while these results suggest that increased use of temporary transfers can lower costs in a single year context, the degree of savings such strategies might produce over the long term is still an open question.

With respect to the solution technique, implicit filtering proves to be an effective search method for the noisy solution (i.e., expected cost) surface generated in this type of water resource problem. The IFFCO algorithm provided repeatable solutions for minimum expected cost and reliability that were accurate to three significant figures. It appears likely that this method may have broader applications within the field of water resource management.

When using a Monte-Carlo based simulation as a component of portfolio optimization “noise” often hinders the optimization process. As simulations grow in size and/or complexity, the value of being able to efficiently manage the noise grows. The control variate method exhibits an ability to reduce the noise such that less than half the number of realizations is required to match the accuracy and precision of the optimal portfolios produced without this method, a result that may lead to this approach finding growing application in water resource planning. In the case of this work, the described approach is applied to the development of multi-year water supply portfolios, allowing for more efficient investigation, as well as an investigation of different long-term option contracts. Some additional creativity in the design of these option contracts, through such techniques as accounting for expected growth in demand or the flexibility to respond to supply conditions, also appears to reduce costs. These portfolio management findings should provide insights useful in future efforts to design water supply strategies in water scarce regions.

Chapter 3: Improved Method for Developing Alternative Hydro-climate Scenarios and their Effects on Water Transfer Agreements

3.1 INTRODUCTION

Traditionally, the ability of a particular water supply to meet demand has been tested against its historic record of precipitation, inflows, and evaporation. Now, in addition to planning for expected demand growths, utilities face the added challenge of anticipating how their water supplies will be affected as climate change drives deviations from historic hydrologic norms.

The International Panel on Climate Change warns of changing precipitation patterns, both in terms of shifting average precipitation values and increased volatility (IPCC 2007). Previous work related to the impacts of climate change on water resource management decisions has largely focused on its effects on existing water supplies (e.g. reservoirs) as opposed to management alternatives such as water transfers (Wood, Lettenmaier et al. 1997; Hamlet and Lettenmaier 1999; Lettenmaier, Wood et al. 1999; VanRheenen, Palmer et al. 2003; VanRheenen, Wood et al. 2004; Vicuna and Dracup 2007). Significant uncertainty remains as to the future effects of climate change, and predictions of how these will impact streamflow profiles (e.g., mean, variance), while currently subject to some skepticism, will continue to improve. As they do, there will be an increased demand for methods of effectively translating predictions regarding changes in the mean and variance into a coherent form representing new streamflow patterns. This work does not aim to improve estimates of streamflow variation due to climate

change, but rather seeks to develop an improved method for producing streamflow records once these estimates have been made, then applies this method within the context of developing water transfer agreements.

In recent years, municipalities have increasingly been turning to transfers as a means of ensuring the reliability of their water supplies (Brown 2006). Numerous studies have examined transfers, their manner of use, and their efficacy in reducing costs while maintaining reliability (Michelsen and Young 1993; Lund and Israel 1995; Lund and Israel 1995; Characklis, Kirsch et al. 2006; Kirsch, Characklis et al. 2009; Palmer and Characklis 2009; Sankarasubramanian, Lall et al. 2009). Nonetheless, little work has evaluated how transfer behavior may be affected by climate change, a potentially important consideration given the long time scales (> 20 years) over which many transfer agreements are often established. This paper evaluates transfer behavior and the performance of several different transfer agreements under different climate scenarios.

Previous research has produced a range of techniques to generate stochastic time series, each with its own advantages and disadvantages. A foundational technique is the bootstrap method, which produces a new time series through a random sampling of the historic record with replacement (Efron 1979). However, this approach produces a time series with no autocorrelation, an important characteristic of hydrologic time series, particularly when considering shorter (e.g. weekly, daily) time steps. Other researchers have developed the bootstrap into a tool that can produce autocorrelated time series via the moving blocks bootstrap (Vogel and Shallcross 1996; Srinivas and Srinivasan 2005) and the nearest neighbor bootstrap (Lall and Sharma 1996). However, the nearest neighbor bootstrap is only able to produce lag-1 autocorrelation (i.e., only the previous

time step has a direct correlative effect upon the current time step). The moving blocks bootstrap is capable of replicating greater levels of the historic autocorrelation as the length of the blocks are extended. This provides for greater autocorrelation, but at the cost of greater replication of the historic record, a potential disadvantage if seeking to explore the impacts of alternative climate scenarios.

A second class of models that have been extensively developed for use in generating streamflow profiles are autoregressive models, which are designed to produce autocorrelated time series. The seminal work was the development of a first-order Markovian autoregressive model (Thomas and Fiering 1962) which led to significant related work (Matalas 1967; Moreau and Pyatt 1970; Jettmar and Young 1975; Young and Jettmar 1976) and culminated with the formalization of autoregressive (AR) models of order p (AR(p)) (Box and Jenkins 1976), part of a class of models that are sometimes referred to as Box-Jenkins models or ARMA models. Autoregressive models, given their ability to accommodate higher-order autocorrelation lags, are particularly effective at replicating the historic autocorrelation. However, AR(p) models assume complete stationarity and impart constant correlation levels for each lag n . Any changes to the rainfall/runoff relationship in a watershed due to seasonal changes in evapotranspiration or infiltration rates that might affect the correlation structure cannot be accounted for by AR(p) models. Neither bootstrap-based methods nor autoregressive techniques are able to both accurately replicate the historic autocorrelation and account for seasonal changes in the correlation structure. The exception is the moving blocks bootstrap, which cannot do both without replicating extended portions of the historic record.

Previous water resource studies have clearly recognized the importance of the presence of autocorrelation in a synthetic time series, and these studies have typically used a monthly timestep. However, as timesteps are shortened, the magnitude of streamflow autocorrelation increases, as does the importance of reproducing those levels of autocorrelation. Failure to reproduce those levels of autocorrelation will prevent the synthetic record from producing long-term extreme events (i.e., droughts) at historic rates.

This work will present a method of generating autocorrelated time series that accurately represent both the seasonal correlation structure and the autocorrelation function. This method, the autocorrelated bootstrap, builds upon the fractional Gaussian noise (FGN) method (McLeod and Hipel 1978). The primary functional difference between the autocorrelated bootstrap and the FGN method is that the development of an FGN model was dependent upon the value of the Hurst statistic, unlike the autocorrelated bootstrap. Moreover, this FGN formulation precludes the consideration of seasonal variations in the correlation structure, and thus produces time series with stationarity in their correlation structure, similar to ARMA models. The primary practical difference between the two methods is the ability of the autocorrelated bootstrap method to produce time series of unlimited length, whereas the FGN method is generally limited to time series of roughly 100 data points or less. (This limitation is not due to less advanced computational capabilities, but rather the sparse nature of the matrices involved.) McLeod and Hipel (1978) used the FGN method on a monthly timestep, and found them to be slightly inferior to autoregressive models according to the Akaike information criterion. However, this was largely the extent of their comparative analysis. There was

no in-depth comparison of autocorrelation or seasonal correlation levels, and their test scenarios involved relatively low levels of autocorrelation. Both the FGN and autocorrelated bootstrap impose a correlation structure upon an uncorrelated time series, a process which requires that the historic data be normalized or “whitened”. When this whitening process is reversed, historic means and standard deviations can be adjusted to reflect changes in the climate, thereby creating a stochastic time series that portrays future climate scenarios while preserving the historic autocorrelation. Results indicate that this approach has the ability to better replicate seasonal variations in the correlation structure and produces streamflow profiles that exhibit rates of extreme flow events more effectively than higher-order autoregressive models.

3.2 METHODOLOGY

3.2.1 *The study region*

This method (fully described in the next section) is applied to a water supply model developed to evaluate a series of risk-based transfer agreements in the Research Triangle region of North Carolina. An earlier version of this model, one using historic streamflow inputs, was described by Palmer and Characklis (2009) and simulates a multi-reservoir system that supplies two utilities, the city of Durham and, via the Orange Water and Sewer Authority (OWASA), the towns of Chapel Hill and Carrboro. Also simulated are the transfer agreements both utilities enter into with a neighboring community (Cary) that has access to a supply with significant excess capacity. Durham operates two primary reservoirs, Lake Michie and Little River Reservoir, with a total capacity of 6.5 billion gallons (BG). OWASA also maintains two primary reservoirs, University Lake and Cane Creek Reservoir, with a combined storage of just under 3.6 BG. The two utilities lie in adjacent watersheds, and streamflows in the two are highly correlated.

Durham's and OWASA's water supplies are expected to be sufficient to meet most demands under most conditions through 2030. However, both utilities recognize the need to augment supplies if they are to meet demand during droughts. They can either develop expensive infrastructure in advance of 2030, an approach likely to result in significant volumes of capacity going unused in the vast majority of years, or they find another manner of augmenting supply on more of an "as needed" basis. Alternatively, Durham and OWASA can enter into a transfer agreement with a third utility, the Town of Cary, which draws its water from Jordan Lake a regional supply with excess capacity, but which Durham and OWASA have no current means to access. Cary would transfer water

via an interconnection with the Durham system, which has an interconnection with OWASA. This is a treated water interconnection, making Cary's treatment plant capacity a potential limiting factor in any agreement (as are the conveyance capacities of the two interconnecting pipes, 11 million gallons per day (MGD) to Durham and seven MGD to OWASA). It is worth noting that transfers of treated water are the most common in the eastern United States, where riparian water law does not generally provide a context (i.e., property rights) in which raw, untreated water can be traded.

Palmer and Characklis (2009) developed a risk-based transfer agreement that triggers the purchase of treated water from Cary whenever the risk of Durham or OWASA's reservoir storage level falling to "failure" (in this case, 20% of capacity) reaches α , a predetermined risk threshold. The transfer agreements are based on risk metrics which evaluate the probability that Durham or OWASA's reservoir storage will fall below its failure level over the next 52 weeks (Figure 3.1), and transfers are requested whenever the specified "risk-of-failure", α , is exceeded. Because treated water is being transferred, the agreements are entered into with the knowledge that Cary must first meet its obligations to its own customers. As a result, Cary can defer some or all of the transfer requests during periods of peak demand. However, these periods are infrequent, and because Durham and OWASA request transfers based on a risk of a future shortfall (i.e., they are not in imminent danger of "running out" of water), this "interruptible" clause in the agreement does not prevent them from reaching their reliability objectives. This interruptible clause allows Cary to defer some summer transfer requests to the fall when it has a higher fraction of unused treatment plant capacity.

The described transfer structure is paired with a hydrologic model in order to evaluate the ability of these risk-based, interruptible transfer agreements. The hydrologic model simulates the reservoirs of Durham and OWASA on a weekly time step using a water balance model. It was not necessary to model the ultimate source of the transfer water, Jordan Lake, given that total withdrawals are expected to be significantly below its safe yield (100 MGD) throughout the simulation period.

The model utilizes 82 years of streamflow and evaporation data, as well as 18 years of demand data (1990-2007). Growth in demand is accounted for by standardizing the weekly demand data, and then scaling up demand in accordance with increases in the mean, projected for future years. In the initial work by Palmer and Characklis (2009), each simulation run was performed for a specific year in the future, meaning that the standardized demand values are all multiplied by the expected average weekly demand value for that future year. The result is that each model run simulated 18 years of reservoir operations assuming the demand levels of a single year in the future.

Converting the hydrologic model to a stochastic form provides the opportunity to change the temporal basis of the simulation. The ability to generate synthetic data beyond the extent of paired historic demand/streamflow data allows the model to be converted to a Monte Carlo simulation, which in turn allows results to be presented in terms of expected values and distributions of likely outcomes. Further, the model is changed so that it no longer simulates a single calendar year repeatedly. Rather, each model realization now simulates a continuous, 16-year period from 2010 to 2025, at a weekly timestep. As a Monte Carlo simulation, the expected outcomes are calculated on the basis of 5000 realizations (i.e., repetitions) of the continuous 16-year period, thereby

providing a comprehensive evaluation of a multi-year transfer agreement. In addition, utilities are concerned not just with expected outcomes, but also extreme events (i.e., high costs). From a collection of model realizations an analysis of high cost events can be obtained.

3.2.2 Managing transfers

The decision to initiate transfers is governed by the risk-of-failure (ROF) metric (i.e., the likelihood that Durham or OWASA's reservoir storage will fall below 20% in the next 52 weeks). Each utility sets a risk threshold, α , such that if the ROF exceeds α , the utility will request transfers and will continue to do so until the ROF falls below α . The ROF values are calculated prior to running the simulation. The ROF is calculated weekly by running a series of 52-week reservoir simulations for each week using each system's initial reservoir storage levels, historic streamflow record, and projected demand values, thereby populating a table that relates the ROF to the calendar week and initial reservoir storage. Figure 3.1 illustrates several risk thresholds (α) for OWASA in the year 2025. In this example, if OWASA sets $\alpha = 2\%$, then transfers will be requested any time the storage level falls below the dashed line, indicating that the ROF is greater than α . Likewise, a risk threshold value of 10% would initiate transfers any time the reservoir level drops below the solid line. Both α lines in Figure 3.1 fluctuate throughout the year. At the beginning of the year, a relatively low reservoir level (e.g., 60%) may not trigger a transfer request because in the winter/early spring there is still an expectation of spring runoff raising storage levels. In the summer, however, that same reservoir level could result in a transfer request, as summer is typically a time of low inflows and reductions in reservoir storage volumes. This work will include consideration of only the 2% and 10%

risk thresholds, values consistent with the range of interest expressed by the utilities and which are perhaps reflective of the risk-averse nature of utilities.




Figure 3.1. Risk chart for OWASA in 2025, given risk threshold values (α) of 2% and 10%.

For the sake of clarity, it should be emphasized that the ROF is purely a historic figure, which is determined using historic hydrologic and demand data. The ROF reports the risk of failure assuming a specific set of conditions (i.e., time of year, reservoir storage volume) and that no other actions are taken by the utility. Initiating transfers would necessarily reduce the likelihood of the water utility suffering a failure in the following 52 weeks. Further, as climate change alters hydrologic conditions, the ROF may no longer be an accurate reflection of the actual risk a utility faces in the absence of

transfers. This modeling effort provides an opportunity to evaluate the effectiveness of anticipatory decision rules based on historic data in a changed environment as well as how these anticipatory decision rules reduce the likelihood of a utility actually experiencing a failure.

Once a request for transfers has been made by either of the buyers (i.e., Durham or OWASA), the seller (Cary) will produce a volume of treated water for transfer that is limited only by a) Cary's excess treatment capacity after meeting its own customers' demands first, less a 5 MGD buffer, and b) the limits of conveyance linking seller to buyer(s). Further, the regional nature of drought implies that OWASA's and Durham's transfer requests are likely to overlap, creating competition for access to the same infrastructure. Thus, a method to apportion transfer requests according to relative levels of risk between competing utilities is used. Clearly, then, requesting water does not necessarily equate to the delivery of water, or even a set quantity of water. Nonetheless, because of the anticipatory nature of the decision rules and the significant level of excess treatment plant capacity in most periods, these agreements still allow the buyers to meet their reliability objectives, assuming suitable risk threshold values (α) are selected. The application of the autocorrelated bootstrap provides an opportunity to test the ability of risk threshold values to avoid water supply failures in a more comprehensive manner than when tested with the more limited historic data set.

3.2.3 Autoregressive models and autocorrelation

Autocorrelation is a special case of the concept of correlation. Whereas correlation measures the extent to which two vectors move in concert, autocorrelation

measures how a time series vector moves relative to itself. The autocorrelation function (ACF) of a time series is expressed via r_k , the autocorrelation at lag k , according to

$$r_k = \frac{1}{N} [\sum_{t=1}^{N-k} (z_t - \bar{z})(z_{t+k} - \bar{z})] / [\sum_{t=1}^{N-k} (z_t - \bar{z})^2]. \quad [3.1]$$

Similarly, a separate analysis can produce the partial autocorrelation function (PACF), which relates the autocorrelation at lag k , independent of lags 1 through $k-1$.

An examination of the ACF and PACF of a time series can provide useful insights into how a time series might be modeled, particularly when significant levels of autocorrelation are present and the modeler is intent upon maintaining those levels of autocorrelation in the synthetic record. For the streamflows considered in this work, high levels of autocorrelation exist in the historic record (Figure 3.5). Water resource studies within the literature commonly utilize a monthly timestep in their models. The shorter (i.e., weekly) timestep in this work leads to higher levels of autocorrelation than what is typically seen with longer timesteps. Likewise, the PACF values are statistically significant at higher lags than what are typically considered in models with monthly timesteps (Figure 3.3), which generally just consider one or two lags of autocorrelation. If a modeler wishes to replicate this level of autocorrelation, examination of the ACF and PACF would indicate that an autoregressive (AR) model may be a suitable choice, as AR models are designed with autocorrelation as a primary goal of the synthetic time series. The general form of the an autoregressive model of order p (AR(p)) is given as

$$z_t = \sum_{i=1}^p \varphi_i \cdot z_{t-i} + \omega_t, \quad [3.2]$$

where ω_t is a white noise process (a random number from a Normal(0,1) distribution) and φ_i is a coefficient such that $|\varphi_i| < 1$. The PACF is used to determine the order of the AR(p) model, in which p is identified as the highest lag for which there is statistically

significant levels of partial autocorrelation. For the case of OWASA, an AR(8) model is appropriate (Figure 3.3).

Despite the ability of AR(p) models to replicate historic autocorrelation levels, several drawbacks exist. The first is that AR(p) models arise out of the study of the ACF and PACF, which evaluate the whole of the time series as a vector. However, streamflows are likely to have variations in their seasonal correlation structure, which are not captured by the ACF or PACF and cannot be replicated by AR(p) models. Second, the time series produced by AR(p) models tend to exhibit increased volatility when systems with higher levels of autocorrelation are simulated. The autocorrelated bootstrap is an attempt to replicate the strengths of autoregressive models while addressing their shortcomings.

3.2.4 *The autocorrelated bootstrap*

Conventional time series analysis often uses the term \vec{y} to refer to vectors of observed values and \vec{z} for synthetically generated values. This work, however, will re-organize the same data into matrices, \mathbf{Y} for historic values and \mathbf{Z} for synthetically generated values. The vectors \vec{y} and \vec{z} are re-formed such that each row of \mathbf{Y} and \mathbf{Z} contain a single year's data. This study utilizes weekly time steps such that \mathbf{Y} and \mathbf{Z} each contain 52 columns. Thus, the (i,j) position of each matrix contains data from the j^{th} week of the i^{th} year.

Developing the autocorrelated bootstrap first requires that the historic data contained in \mathbf{Y} must be seasonally de-trended and made to approximate a Normal(0,1)

distribution, or “whitened”. In the case of inflows, the first step is to take the log of the inflows in \mathbf{Y} to produce $\hat{\mathbf{Y}}$, which is then whitened according to

$$Y_{i,j} = (\hat{Y}_{i,j} - \bar{Y}_j) / \sigma_{\hat{Y}_j} \quad [3.3]$$

where \bar{Y}_j is the mean of column j of $\hat{\mathbf{Y}}$, and $\sigma_{\hat{Y}_j}$ is the standard deviation of column j of $\hat{\mathbf{Y}}$.

It is this white data \mathbf{Y} , and later \mathbf{Z} , that are converted back into \vec{y} and \vec{z} to determine their respective ACF and PACF.

To generate a synthetic correlated time series, the process begins with a bootstrap from the historic flow matrix \mathbf{Y} to populate a matrix \mathbf{X} , which now contains uncorrelated flow data. Specifically, the creation of \mathbf{X} is arrived at via an intermediate matrix \mathbf{M} . Assuming that the historic record contains N years (\mathbf{Y} is therefore a $N \times 52$ matrix), an intermediate matrix \mathbf{M} is formed such that, for each i and j , the value of $\mathbf{M}_{i,j}$ is sampled with replacement from the set $[1, 2, \dots, N]$. A matrix of uncorrelated flow values \mathbf{X} is then formed so that $X_{i,j} = Y_{(\mathbf{M}_{i,j}),j}$, where the number of rows of \mathbf{X} and \mathbf{M} is equal to the number of years to be modeled. While matrix \mathbf{M} plays a minor role in this bootstrap, it will later be seen to be a key component in maintaining cross-correlation between different time series.

To convert the uncorrelated time series to a correlated one, the correlation matrix of the historic record, $\mathbf{Corr}(\mathbf{Y})$, is calculated. The (i,j) position of $\mathbf{Corr}(\mathbf{Y})$ corresponds to the correlation of column i to column j in \mathbf{Y} . A Cholesky decomposition can be performed on $\mathbf{Corr}(\mathbf{Y})$ such that

$$\mathbf{Corr}(\mathbf{Y}) = \mathbf{Q}\mathbf{Q}^T \quad [3.4]$$

in which \mathbf{Q} and its transpose \mathbf{Q}^T are upper and lower triangular matrices, respectively.

The matrix \mathbf{Q} imposes the historic correlation structure upon the uncorrelated matrix \mathbf{X} according to

$$\mathbf{Z} = \mathbf{XQ}, \quad [3.5]$$

such that, on average, $\mathbf{Corr}(\mathbf{Z}) = \mathbf{Corr}(\mathbf{Y})$. Given the upper-triangular nature of \mathbf{Q} , the value of each element $\mathbf{Z}_{i,j}$ is comprised of contributions from the product of each uncorrelated flow value $\mathbf{X}_{i,k}$ (where $k \leq j$) and $\mathbf{Q}_{m,j}$ (where $m \leq i$). It is in this formulation that [3.5] is mathematically analogous to an autoregressive model.

Even though the correlation matrix of \mathbf{Y} has been imposed on the matrix \mathbf{Z} , when \mathbf{Z} is converted into a vector \vec{z} , its autocorrelation function (ACF) indicates less autocorrelation than exists in the historic vector \vec{y} . Any correlation information contained within \mathbf{Q} originates from within the $\mathbf{Corr}(\mathbf{Y})$ matrix, which contains no inter-annual correlation information (i.e., $\mathbf{Corr}(\mathbf{Y})$ contains the correlation of week 1 to week 52 of the same year, but none regarding the correlation of week 52 to week 1 of the following year). Thus, \mathbf{Z} is only intra-annually correlated, and the autocorrelation of \vec{z} is disjoint every 52 weeks.

The disjointed nature of the autocorrelation of \vec{z} can be addressed by employing a matrix manipulation technique, a visualization of which is shown in Figure 3.2A. This technique takes the historic flow data contained within \mathbf{Y} and re-organizes it into a new matrix \mathbf{Y}' that contains inter-annual correlation information. The correlation information contained within both \mathbf{Y} and \mathbf{Y}' will be applied to an uncorrelated time series contained within the matrix \mathbf{X} and will ultimately result in a synthetic time series \vec{z} with similar autocorrelation to that of the historic vector \vec{y} . The matrix \mathbf{Y} is re-formed into vector \vec{y} ,

and the first and last 26 weeks are removed from \vec{y} , which will now be denoted as \vec{y}' . The shortened vector \vec{y}' is re-formed into the matrix \mathbf{Y}' . This new matrix now contains one fewer row than \mathbf{Y} , having had 52 weeks trimmed from the data set. Each column of \mathbf{Y}' still contains a specific week's observations. However, the first column of \mathbf{Y}' contains data from week 27, and the fifty-second column contains data from week 26 of the following year. Now $\mathbf{Corr}(\mathbf{Y}')$ contains correlation data linking the last half and first half of consecutive years, and is used to calculate \mathbf{Q}' as in [3.4]. This conversion method assumes that the historic record is sufficiently long that the loss of one year's data will not significantly alter correlation values between two weeks in the same 26-week block. That is, the correlation between any two weeks within weeks 1 to 26 (or within weeks 27 to 52) should be the same in both $\mathbf{Corr}(\mathbf{Y})$ and $\mathbf{Corr}(\mathbf{Y}')$.

The process of incorporating consideration of both the inter- and intra-annual correlations into the development of the synthetic flow record begins by bootstrapping an uncorrelated flow matrix, \mathbf{X} . The matrix \mathbf{X} is converted to \mathbf{X}' just as \mathbf{Y} was converted to \mathbf{Y}' , and similarly, \mathbf{X} loses one year of data in the conversion to form \mathbf{X}' . Therefore, \mathbf{X} must be formed with one more year's worth of data (i.e., one more row) than the modeler intends to simulate. Equation 3.5 is performed with \mathbf{X} and \mathbf{Q} , and again with \mathbf{X}' and \mathbf{Q}' , to produce the correlated inflows \mathbf{Z} and \mathbf{Z}' , respectively (Figure 3.2B).

The matrices \mathbf{Z} and \mathbf{Z}' are alike in that the correlation contained within each has been imposed in 52-week segments (via [3.5]), but these segments begin and end at two different points in the calendar year. However, because \mathbf{Z} and \mathbf{Z}' both originate with data contained within the same uncorrelated flow vector \vec{x} , data from the two matrices can be merged to form a combined matrix \mathbf{Z}^C , which, when converted into vector \vec{z}_c ,

produces a time series that is inter-annually correlated and lacks the disjoint nature of .

Figure 3.2B illustrates which 26-week segments are selected from \mathbf{Z} and \mathbf{Z}' and how they are composed to form the matrix \mathbf{Z}^C . Note that the 26-week segments that comprise \mathbf{Z}^C originate from the right halves of \mathbf{Z} and \mathbf{Z}' , which are correlated to the 26-week segments on the left-hand sides of \mathbf{Z} and \mathbf{Z}' . More plainly, each datum point in \mathbf{Z}^C contains correlation information from a minimum of 26 previous data points in the time series.

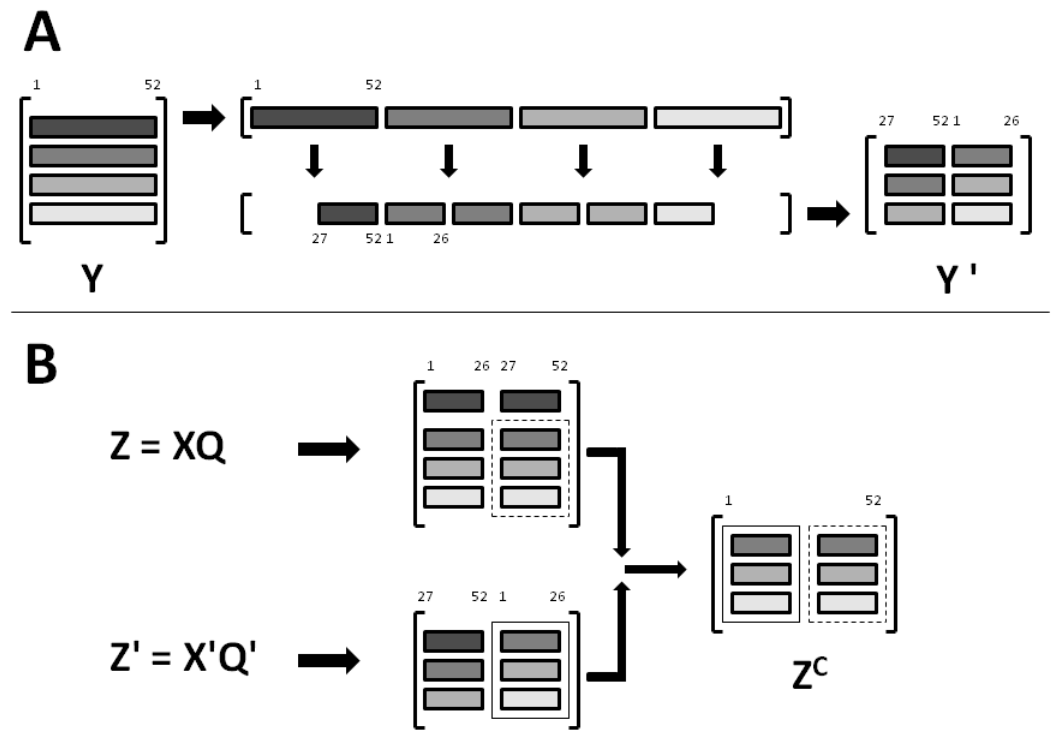


Figure 3.2. Part A illustrates how \mathbf{Y} is converted to \mathbf{Y}' (as well as how \mathbf{X} converts to \mathbf{X}'). Part B illustrates how \mathbf{Z} and \mathbf{Z}' are combined to form \mathbf{Z}^C .

Figure 3.3 shows the partial autocorrelation function (PACF) of the 82-year historic inflows to one of the water systems in the study region (the Orange Water and Sewer Authority, or OWASA) as well as that of an 82-year synthetic record created using




Figure 3.3. Partial autocorrelation function of historic and stochastic OWASA inflows.

the autocorrelated bootstrap. The historic PACF indicates statistically significant partial autocorrelation to about lag eight. The manner in which each datum point in \vec{z}_c incorporates correlation information from at least 26 previous data points is analogous to the manner in which autoregressive models incorporate correlation information from previous data points in the time series. Thus, each \mathbf{Z}_{ij}^C value is constructed in a manner similar to an autoregressive (AR(p)) model, where $p \geq 26$, but the PACF shows that on average only the eight most recent uncorrelated flow values contribute significantly to the value of \mathbf{Z}_{ij}^C . However, the correlation structure can change seasonally, and Figure 3.4 illustrates how the historic one-lag correlation fluctuates throughout the year for OWASA. This is compared to the one-lag correlation produced by the autocorrelated bootstrap method and an AR(8) model, and the inability of the AR(8) model to adjust for seasonal correlation levels is clear. Figure 3.5 demonstrates the effectiveness of the

autocorrelated bootstrap's ability to replicate the historic ACF, again using the inflows to OWASA's reservoirs as a point of comparison.

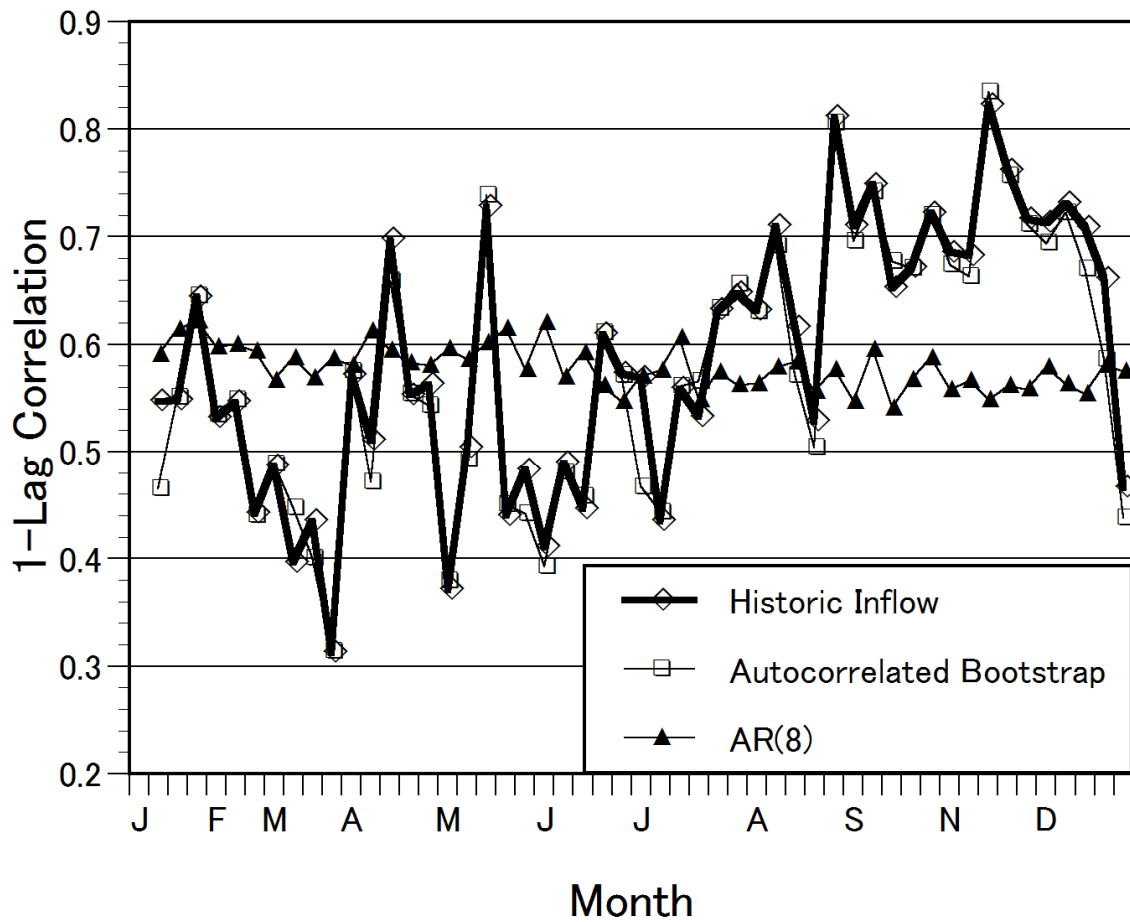


Figure 3.4. The 1-lag correlation of historic OWASA inflows, obtained from $\text{Corr}(Y)$, and compared to those produced by the average values of a sample of time series produced by an autocorrelated bootstrap and an AR(8) model.

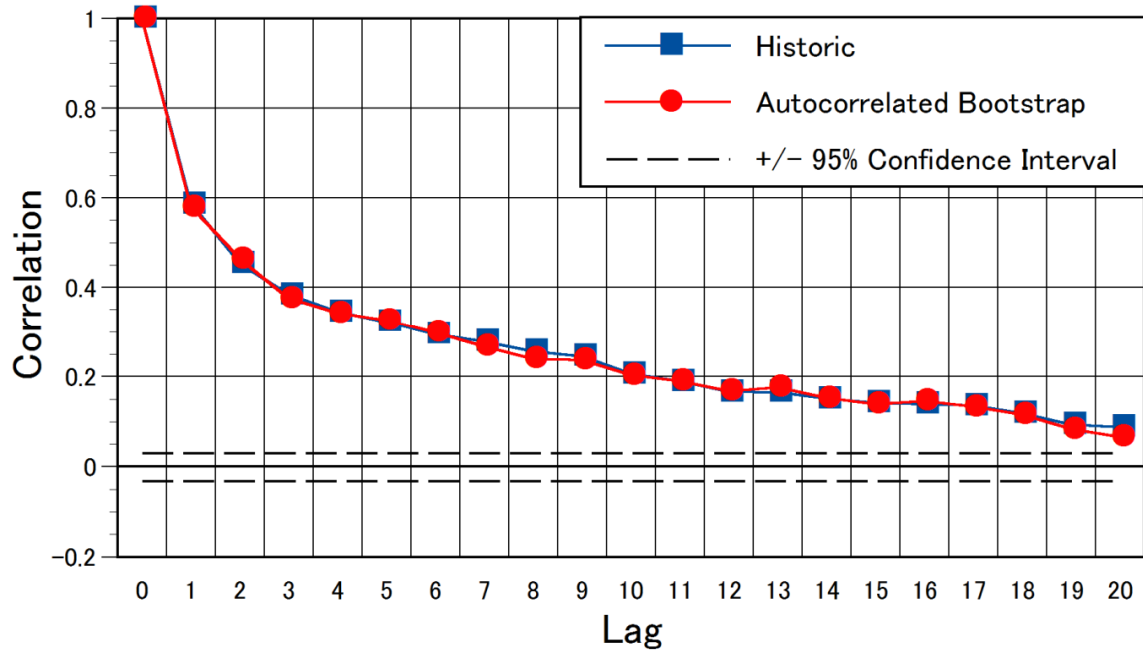


Figure 3.5. Autocorrelation function of historic and stochastic OWASA inflows.




Figure 3.6. Histograms of Standardized Inflow Index values calculated for OWASA inflows, produced from the historic time series and the mean values of a sample of uncorrelated bootstraps and autocorrelated bootstraps, both bracketed by their 25th and 75th percentiles.

A separate evaluation of the autocorrelated bootstrap approach involves assessing whether more extreme events (e.g., droughts) are produced with the same frequency and severity as in the historic record. For this, the Standardized Precipitation Index (SPI) (McKee, Doesken et al. 1993; Guttman 1999) will be adapted for use with inflows. The SPI is a drought index that uses a running total of precipitation across a user-defined interval (commonly one, three, or six months) to produce a z-score, whereby a score of +1.0 indicates that precipitation is one standard deviation above normal. At each time step in the historic record, an index is created relating the cumulative flow to the SPI value (i.e., z-score). The SPI is a methodology that analyzes seasonally fluctuating time

series, and is therefore useful in this context, but since it is being adapted for use with streamflows, it will be referred to here as the Standardized Inflow Index (SII). Figure 3.6 displays a histogram of the historic three-month SII values, the mean histogram of a sample of autocorrelated bootstraps, and the mean histogram of a time series formed from a plain bootstrap. Both histograms of synthetic time series are bracketed by their twenty-fifth and seventy-fifth percentiles (light dotted lines). The autocorrelated bootstrap shows good agreement with the historic SII histogram, particularly at high and low values. Note that in the absence of autocorrelation, deviations from the mean are significantly reduced. This greatly reduces the frequency of major droughts and high flow events, which in turn serves to moderate transfer behavior. An eighth-order autoregressive model was created and performed similarly to the autocorrelated bootstrap. However, the AR(8) model significantly overproduces high flow events ($SII > 2$) (Figure 3.7). Overestimation of high flow events is problematic not just in that it represents a deviation from historic flow patterns, but also that those high flows may be stored in reservoirs, thereby moderating the effects of future low flow events.




Figure 3.7. Histogram of SII values for high flow events (> 2) for OWASA, produced from historic data and the mean results of 20 autocorrelated bootstraps and AR(8) time series, both bounded by their 25th and 75th percentiles.

3.2.5 Application to the study region

The autocorrelated bootstrap is applied to the study region to generate both reservoir inflows and evaporation values. Evaporation is used within the reservoir model to account for the net flux of water through the reservoir surface area, and the historic evaporation record extends across the same 82 years as the historic inflow record. There exists a significant amount of crosscorrelation between OWASA's and Durham's inflows (correlation coefficient of 0.95), and, to a lesser extent, between inflows and evaporation.

In order to retain some level of crosscorrelation, only a single matrix \mathbf{M} is developed (in which each $\mathbf{M}_{i,j}$ is sampled with replacement from the set $[1,2,\dots,N]$) which is then used to create individual uncorrelated flow matrices \mathbf{X} for both Durham and OWASA, as well as uncorrelated evaporation data. The fact that all three \mathbf{X} matrices originate from the same matrix \mathbf{M} means that, while the time series contained within them are not autocorrelated, they are crosscorrelated. The use of a single \mathbf{M} matrix preserves the majority of crosscorrelation, with Durham and OWASA's stochastic inflows' crosscorrelation coefficient averaging 0.89.

The demand data lack a significant level of both autocorrelation and correlation to inflows, and as such the stochastic hydrologic model generates demand records through an uncorrelated bootstrap of the standardized demand data. For each model realization, the 16 years of stochastic demand data that is generated is un-standardized, year by year, such that the first 52 weeks of demand data is multiplied by the projected average weekly demand in 2010, the second 52 weeks by expected demand in 2011, and so on.

The final step of the autocorrelated bootstrap is to reverse the whitening process described in [3.3] in which the historic data was made to approximate a Normal(0,1) distribution. This “un-whitening” process reverts the data contained in \mathbf{Z}^C back to the original distribution of the historic data, with each column j of \mathbf{Z}^C multiplied by $\sigma_{\hat{Y}_j}$ before the mean \bar{Y}_j is added. This process presents an opportunity to adjust the autocorrelated time series to reflect future climate change scenarios by slightly altering the values of $\sigma_{\hat{Y}_j}$ and \bar{Y}_j .

Four climate scenarios are considered within the study region. These scenarios are selected not as a result of specific predictions of future conditions, but rather as a

reasonable range of potential hydrologic changes that can be used to evaluate how such representative changes might impact transfer behavior. The first scenario is based on the historic record, forming a base case which reflects the historic statistical conditions and is labeled “Hist.”. The IPCC (2007) warns, without citing specific figures, of increased precipitation volatility as a likely outcome of climate change. Consequently, the second and third scenarios examine scenarios representative of increases in the standard deviation of inflows of 10% ($\sigma_{\hat{Y}_j} \cdot 1.1$) and 20% ($\sigma_{\hat{Y}_j} \cdot 1.2$), labeled as “SD 10” and “SD 20”, respectively. The fourth scenario reproduces the statistical conditions of the most recent decade “MR 10”, a particularly dry period that some have posited as being representative of future conditions. This period includes two severe droughts, and involves mean annual inflows 7% lower than in the full historic record and weekly standard deviations in inflows that average 12% greater.

3.3 RESULTS

The results presented here represent not just the four climate scenarios, but also three transfer levels: no transfers, transfers with a risk threshold (α) of 2%, and 10%. It should also be stressed that the hydrologic statistics remain constant across the 16-year simulation period, and that any change in transfer behavior over time is a reflection of water growth in demand within the Durham and OWASA service areas.

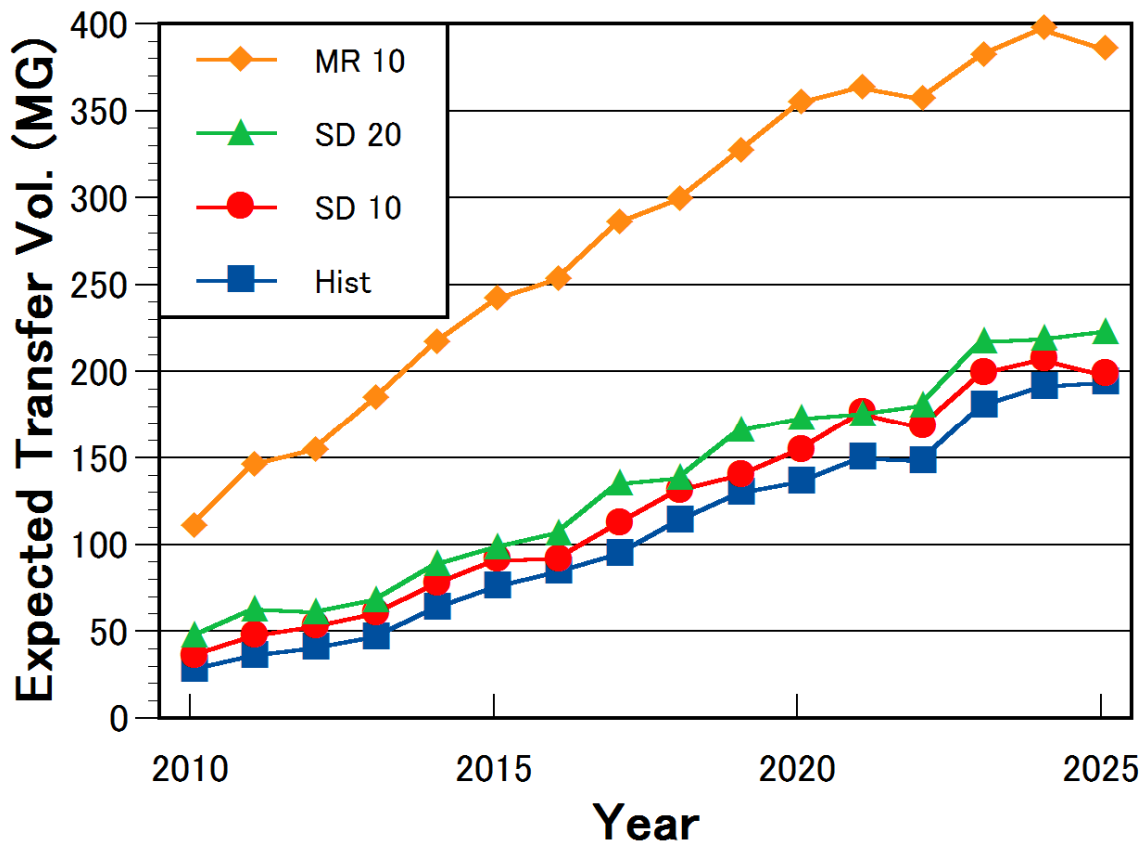


Figure 3.8. Expected annual volume of water transferred to Durham, given $\alpha = 2\%$.

Examining directly the effect the four climate scenarios have on transfer behavior, Figure 3.8 presents the expected annual volume of transfers for Durham over time, using a risk threshold α of 2%. Unsurprisingly, the volume of water transferred increases over

time, but notably, increasing the volatility of the inflows (SD 10 and SD 20) shifts the trendline upward relative to the base case (Hist.). This indicates that increased volatility alone serves as a constant “fee” in terms of an increase in expected transfers, regardless of the base rate of expected transfers. The MR 10 scenario presents a marked departure from the other three scenarios and demonstrates the extreme sensitivity of the system to even slightly reduced inflows. The reduced inflows cause a significant shift upward, as well as a steeper slope to the expected transfer growth rate, roughly doubling the expected transfers in the MR 10 scenario, relative to the other three.

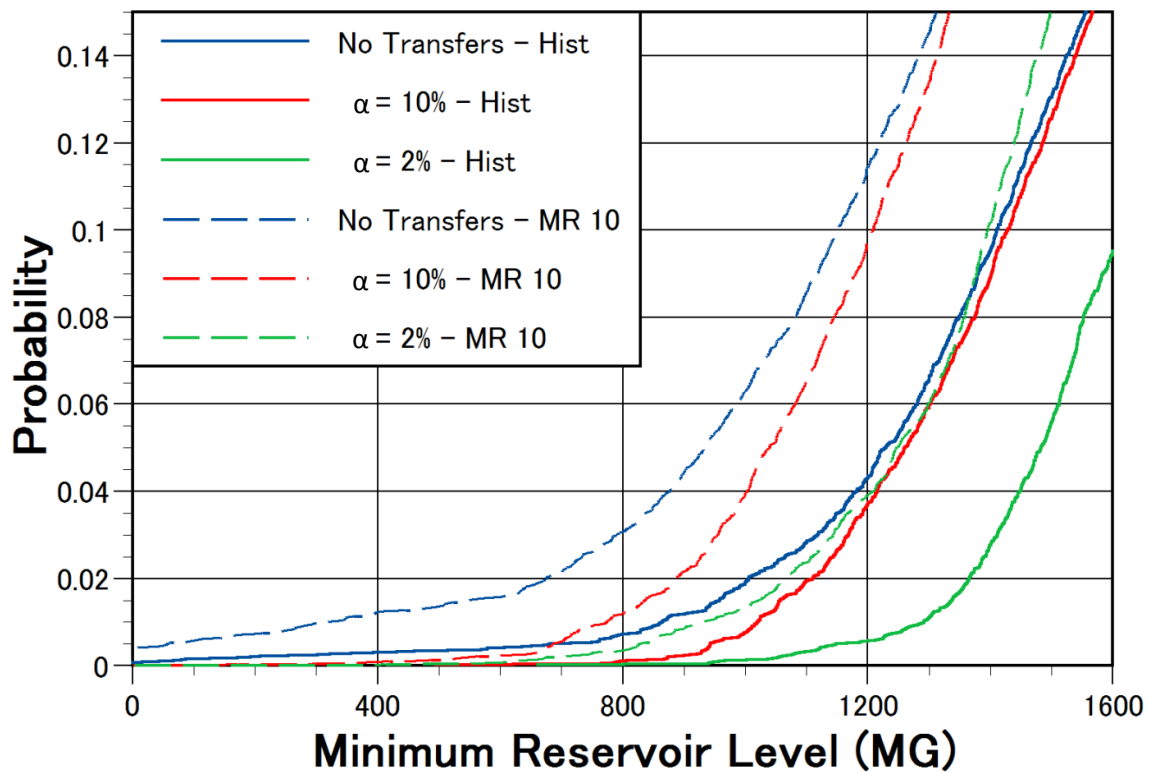


Figure 3.9. Cumulative distribution functions of the lowest observed annual storage values for OWASA in 2025.

Model results (involving 5000 realizations) are used to produce cumulative distribution functions (CDF) of the lowest observed reservoir levels for OWASA in 2025 (Figure 3.9). Six CDF curves are presented, representing the three transfer levels (no transfers, $\alpha = 2\%$, and $\alpha = 10\%$) under both the historic and MR 10 climate scenarios. The vertical line at 710 MG represents the failure level for OWASA's reservoirs (20% of available storage), and it can be seen that without transfers, the MR 10 scenario represents roughly a tripling of the likelihood of a failure occurring in 2025 (from 0.7% to 2.2%) compared to the historic climate scenario. Further, given the MR 10 scenario, using higher a risk threshold value (α) of 10%, the likelihood of a failure is only reduced to the same level as that of the historic climate scenario without the aid of transfers. For all three transfer scenarios, shifting from the historic climate to the MR 10 scenario shifts the CDF curves towards greater risk and higher likelihood of lower reservoir volumes.

A similar analysis of CDF curves is performed for Durham for Figure 3.10. Here, the MR 10 scenario is analyzed in 2025 with no transfers and at α values of 2% and 10%. The vertical line at 1270 MG represents Durham's failure level. Left of this failure line, the CDF curves of the two transfer scenarios are nearly identical. If the goal of employing transfers is to avoid failures, then the 2% risk threshold is only marginally better than the 10% risk threshold, but at a much higher cost. The expected annual cost of transfers corresponding to the 10% risk threshold in 2025 is \$433,000, compared to an expected annual cost of \$963,000 at $\alpha = 2\%$, an increase of 123%. That said, above the failure point the two different transfer regimes differentiate themselves, with the likelihood of reservoir storage falling to 2500 MG in a given year (twice the volume of

the failure point) approximately 40% lower engaging a risk threshold of 2% as opposed to that of 10%.

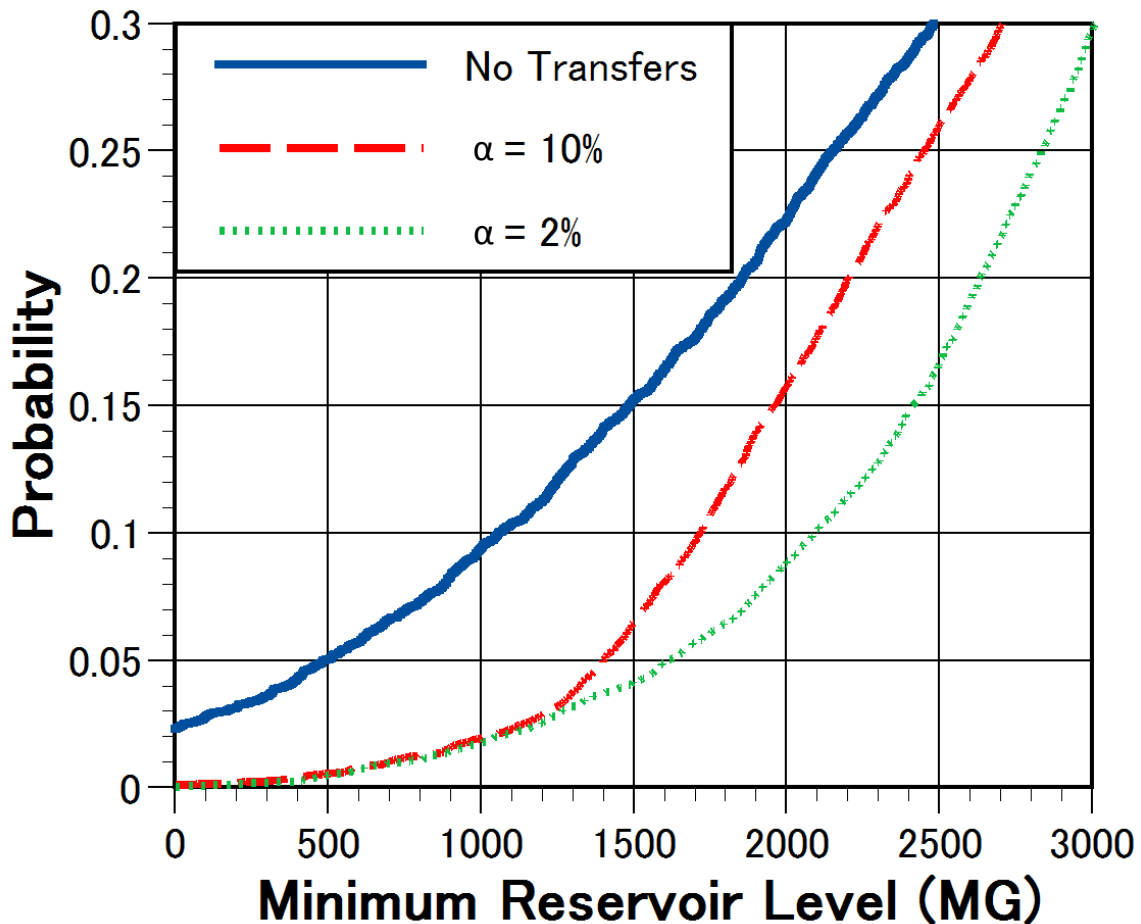


Figure 3.10. Cumulative distribution functions of the lowest annual reservoir storage values for Durham in 2025 for the MR 10 scenario.

Three basic factors determine the shape and position of these CDF curves: climate, infrastructure (i.e. the rate at which water can be transferred), and policy (i.e. the risk threshold values). Figure 3.9 presents the effect of climate and its impact on the risk profile to a utility, as well as how different policy choices can mitigate those risks. In

Figure 3.10, to the right of the failure point, the differentiation of the 2% and 10% risk threshold CDF curves is policy driven, as both scenarios share the same climate conditions and infrastructure. However, the convergence of those two CDF curves left of the failure point indicate that, for this range of risk threshold values, the system is constrained by infrastructure. Significant improvement to the risk profile in that region would require either an increase in infrastructure capacity or a lowering of the risk threshold value to initiate transfers at an earlier point (i.e., initiate transfers while reservoir levels are higher). It should be noted that these results reflect the year (2025) with the highest demand of any year considered, yet both risk thresholds provide an annual reliability of approximately 97%.

Transfers reduce the excess capacity that a utility must constantly maintain, but the selection of certain decision variables can lead to inefficiencies as well. Selection of a low risk threshold can lead to greatly increased transfer rates with little or no advantage in avoiding failures compared to some higher risk thresholds, as the decision variable requires action before a crisis that may or may not materialize. To continue with the example of Durham in 2025, without transfers under the historic climate scenario, failures occur in 118 years out of the 5000 annual observations for 2025 (a failure rate of 2.4%). However, given $\alpha = 10\%$, transfers occur in 800 years out of 5000 observations (16% annual rate), and when α is reduced to 2%, the number of years with transfers jumps to 3053 (61% annual rate) for the very same climatic time series without a significant difference in failure rates (both reduce the annual failure rate by a factor of 10).

Table 3.1. Total costs of transfers in millions of dollars, 2010 – 2025.

	OWASA Hist.	OWASA MR 10	OWASA Hist.	OWASA MR 10	Durham Hist.	Durham MR 10	Durham Hist.	Durham MR 10
α	2%	2%	10%	10%	2%	2%	10%	10%
Average Cost (NPC)*	\$0.90 (\$0.48)	\$1.55 (\$0.85)	\$0.13 (\$0.07)	\$0.47 (\$0.26)	\$4.26 (\$2.60)	\$11.13 (\$7.09)	\$1.40 (\$0.83)	\$5.69 (\$3.50)
Median Cost (NPC)	\$0.65 (\$0.34)	\$1.09 (\$0.58)	\$0 (\$0)	\$0 (\$0)	\$3.96 (\$2.37)	\$10.93 (\$6.91)	\$0.76 (\$0.44)	\$5.23 (\$3.20)
Conditional- Value-at-Risk (NPC)	\$3.92 (\$2.14)	\$5.99 (\$3.38)	\$2.12 (\$1.16)	\$4.46 (\$2.50)	\$10.14 (\$6.37)	\$20.16 (\$13.26)	\$6.23 (\$3.75)	\$14.40 (\$9.09)
Likelihood of No Transfers	13.4%	2.6%	86.4%	64.6%	0.2%	0.0%	32.6%	2.2%
Average Time to First Transfer (in Years)**	11.35	10.52	13.13	12.23	5.74	2.76	9.77	6.09

*2010 Net Present Costs are calculated with a 5% discount rate

**Excluding realizations in which no transfers occur

The cumulative costs of the transfer program over the 2010 - 2025 study period are presented in Table 3.1 for a 2% risk threshold. The costs of the transfer scenarios are based on a sales price of \$2.50 per thousand gallons (kgal) to Cary, and for OWASA, a wheeling fee of \$0.50/kgal that must be paid to Durham for passing transferred water from Cary to OWASA through its distribution system. Revenues from the wheeling fee are not deducted from Durham's costs, but would result in a net lowering. Where costs are presented in Table 3.1, the parenthetical figures below them represent the net present cost, assuming a 5% discount rate.

The cost differentials between OWASA and Durham are obvious, reflecting both OWASA's greater drought resilience and Durham's larger size. It should be noted that a

number of model realizations never result in transfers during the course of the 16 year simulation, and thus no transfer costs are incurred. The likelihood of zero transfer costs is reflected in the table, with OWASA's likelihood dropping from 13.4% for the historic scenario down to 2.6% for the MR 10 scenario. The likelihood of not transferring water between 2010 and 2025 is much greater using a 10% risk threshold (not included in Table 3.1), ranging from 86.4% to 64.6% for the same two climate scenarios for OWASA. This reflects some of the costs savings available with a greater risk tolerance, as results suggest a median cost of \$0 for all four climate scenarios at the $\alpha = 10\%$ level. For both Durham and OWASA, a 10% increase in the standard deviations of reservoir inflows causes roughly a 10% increase in average and median costs. The rise in transfer activity observed between the historic and MR 10 climate scenarios results in significantly larger increases in costs. Durham's expected costs increase by over 160%, while OWASA's increase by over 70%. If the future climate is similar to that of the previous decade, the expected cost increase in terms of OWASA and Durham's transfers is \$7.5 million, an increase of 146% relative to the historic record. Table 3.1 also lists the average time to the first transfer in a given model realization (assuming a transfer occurs in that realization) given a start date of 2010. While the expected time to first transfer drops by slightly less than one year for OWASA from the historic climate scenario to MR 10, for Durham the time to first transfer drops by three years, information that could lead to an acceleration of any supply augmentation plans.

Average and median costs are clearly of interest to utilities, but as the presence of realizations with no transfers (and thus no costs) demonstrate, metrics of cost volatility are vital as well. The Conditional Value at Risk (CVAR) provides a metric for the likely

worst-case scenario by reporting the expected cost, assuming the exceedance of the ninety-fifth percentile of costs (i.e., the average of all costs between the ninety-fifth and one-hundredth percentiles). On a relative basis, OWASA has greater cost volatility concerns than Durham, with OWASA's CVAR values representing roughly a four-fold increase over expected costs. Durham's CVAR values, on the other hand, are approximately double that of its expected costs. Durham's lower relative cost volatility is a result of its more consistent reliance on transfers.

3.4 CONCLUSIONS

Assessing the behavior of water supply policies could be done using the historic record, but even a lengthy historic record may only provide a handful of significant droughts. Stochastic modeling can produce large data sets that allow the full range of likely outcomes to be evaluated. Further, the ability to simulate the effects of a changed climate is especially important given the long-term planning horizon utilities rely upon. Significant uncertainty remains as to the effects of climate change, particularly as they relate to hydrologic impacts. While these predictions are still being refined, the development of a stochastic method that produces streamflow inputs consistent with historic statistical properties, and which can also be adjusted to reflect future hydrologic conditions, is a step forward. Results indicate that by being able to recreate the seasonal correlation structure, the autocorrelated bootstrap better simulates historic distributions, which in turn suggest that estimates of climate change impacted streamflows are more consistent as well.

The autocorrelated bootstrap was shown to successfully replicate both the autocorrelation and the partial autocorrelation record while more accurately producing extreme flow events at a rate observed in the historic record. Unlike autoregressive models, the autocorrelated bootstrap is capable of recreating the seasonal correlation structure of the historic record. Further, the un-whitening process provides an opportunity to adjust the basic statistical properties of the inflow record to evaluate the effect of a range climate change scenarios.

Utilities planning for future demand are increasingly considering creative supply alternatives, such as temporary transfers. As opposed to the up-front capital costs of

traditional water supply expansion projects (e.g. reservoirs), transfers used on an “as-needed” basis require utilities to consider significant intermittent and variable costs.

Also, the utilities must be assured that the policies that control transfer use provide the necessary levels of reliability in an efficient manner.

The results show that despite using a risk-based decision rule developed using the historic record, transfers were still sufficient to prevent failures in all but the most severe droughts, even under the more challenging climate scenarios. Results also demonstrate that increased volatility in a water supply’s inflows causes an increase in transfers, and thus, cost. More telling is that costs are extremely sensitive to a change in expected inflows. A small reduction in mean inflows (7%) produces a two-fold increase in costs. If the next 16 years in the Research Triangle region are similar to the last decade hydrologically, the expected costs of the transfer program described will be substantially more expensive for local utilities than if the climate is similar to the last 82 years. This work highlights the importance of developing accurate climate change models and determining how climate change impacts can affect utilities’ long-term planning efforts.

Chapter 4: Conclusions

Growing cities and towns are challenged to maintain water supply reliability as their demands for water increase. The cost and difficulty of maintaining the volume of firm capacity required to meet reliability goals is leading many utilities to consider alternatives, such as water transfers. This work reinforces the findings that transfers can be a more cost-effective means of ensuring reliability than firm capacity expansion. Transfers, though, require a higher level of management than typical firm capacity, and the tools developed in this work can provide utilities confidence that they are employing transfers in an efficient manner.

Addressing water resources with the use of stochastic modeling provides several advantages. One is that, absent omniscience, the use of expected outcomes is the most appropriate metric by which to evaluate water supply alternatives. Further, particularly through the use of Monte Carlo techniques, the range of likely outcomes can be explored as well. However, Monte Carlo methods have been of limited use to water resource planners due to the difficulty of optimizing such models. The work in Chapter 2 addressed this challenge by pairing a Monte Carlo model with an optimization routine known as implicit filtering, which is capable optimizing “noisy” solution surfaces that are characterized with high frequency, low amplitude perturbations. This simulation-optimization pairing represented the first time that a Monte Carlo model has been optimized within the water resources field.

When presented with a variety of market-based water supply assets, it is seen that the assets can be combined to form a portfolio, and, taking advantage of the properties of

individual assets, the composition of assets can be adjusted to achieve particular properties (i.e., reliability, cost volatility). The results from the LRGV indicate that there is a general tradeoff of lower costs in return for higher cost volatility. If a utility has reservations about relying on the spot market due to price volatility, options may present a viable alternative. While options are priced to be risk-neutral relative to spot market leases, in practice some fraction of options lapse, even in optimized portfolios. The average unit cost of water obtained through the exercise of options, then, is higher than the average unit cost of leases purchased in the option exercise month, the difference in unit costs being dependent upon the average fraction of lapsed options. The difference in average costs acts, in essence, as a form of insurance against price spikes in the spot market. A utility is unlikely to desire to renegotiate option contracts on an annual basis. As such, long-term modeling efforts can inform as to the performance of long-term option contracts. Here, a 10-year long option contract was modeled as being representative of the timescale in which a utility is likely to be interested. On that timescale, issues of growth in demand must be considered by the utility. This work showed that with slight alterations to the option contract in how it addresses demand growth, significant savings can be achieved.

The expansion of the Lower Rio Grande Valley (LRGV) model from a one-year to a 10-year model shows the rapidity with which the computational burden can expand. The success of the control variate (CV) method in reducing model variance was significant, reducing by over half the number of model realizations required to achieve a given level of accuracy and precision in the optimized results. The CV method relies upon the modeler's familiarity with the sources of variability within a system. Many

water resource simulations are custom-designed for specific water sheds, and therefore are likely to have one or more modelers that may have the requisite knowledge of variance within the system to implement the CV method. Demonstrating the ability of Monte Carlo simulations to be optimized using implicit filtering may lead to the wider adoption of this methodology within the water resources community, and with it, the potential for the CV method to be utilized.

The modeling effort in the Triangle (Chapter 3) reinforces many of the lessons learned about transfers in the LRGV, particularly regarding the efficiency and cost-effectiveness of ensuring water supply reliability via transfers. Month-long timesteps are often used in water resource models, and a reduction to a weekly timestep increases the level of autocorrelation that is likely to be encountered. Reproducing the level of autocorrelation seen in the historic streamflow record is vital if droughts are to be simulated at rates commensurate with the historic record. The autocorrelated bootstrap (AB) method was proven to be adept at reproducing the historic auto- and partial autocorrelation functions. Unique to the AB method is the ability to recreate the auto- and partial autocorrelation functions as well as the seasonal correlation fluctuations. Other stochastic time series generation methods assume that the lag n correlation is constant throughout the seasons. In reality, seasonal fluctuations in evapotranspiration rates and soil moisture can affect the lag n correlation. Moreover, the AB method can be implemented in a manner to maintain crosscorrelations between time series. Similar to the LRGV model, the Triangle model is a long-term Monte Carlo simulation that is able to evaluate the effect of growing demands on transfer behavior and effectiveness. In addition, the “whitening” process involved with the AB method provides an opportunity

to evaluate transfer behavior under different climate scenarios. As climate change modeling improves its predictive ability with regards to precipitation, the AB method can act as a powerful tool to translate those predictions to actual effects on water resources.

Taken together, the tools developed in this work can provide utilities with the confidence to implement transfer strategies as a lower cost alternative to firm capacity expansion.

References

- Anderson, T. L. and P. J. Hill, Eds. (1997). Water Marketing - The Next Generation. Lanham, MD, Rowman & Littlefield.
- Audet, C. and J. E. Dennis (2003). "Analysis of generalized search patterns." SIAM J. Optim. **13**: 889-903.
- Avramidis, A. N. and J. R. Wilson (1996). "Integrated Variance Reduction Strategies for Simulation." Operations Research **44**(2): 327-346.
- Batterman, A., J. M. Gablonsky, A. Patrick, C. T. Kelley, T. Coffey, K. Kavanagh and C. T. Miller (2002). "Solution of a groundwater control problem with implicit filtering." Optim. Eng. **3**(189-199).
- Black, F. and M. Scholes (1973). "The pricing of options and corporate liabilities." J. Polit. Econ. **81**: 637-659.
- Black, F. and M. Scholes (1973). "The Pricing of Options and Corporate Liabilities." Journal of Political Economy **81**: 637-659.
- Box, G. E. P. and G. M. Jenkins (1976). Time Series Analysis: forecasting and control. San Francisco, Holden-Day.
- Boyle, P. P. (1977). "Options: A Monte Carlo Approach." Journal of Financial Economics **4**(3): 323-338.
- Broadie, M. and P. Glasserman (1996). "Estimating Security Price Derivatives Using Simulation." Management Science **42**(2): 269-285.
- Brookshire, D. S., B. Colby, M. Ewers and P. T. Ganderton (2004). "Market Prices for Water in the Semiarid West of the United States." Water Resources Research **40**.
- Brown, T. C. (2006). "Trends in water market activity and price in the western United States." Water Resources Research **42**(9).
- Chang, C. and R. C. Griffin (1992). "Water Marketing as a Reallocative Institution in Texas." Water Resources Research **28**(3): 879-890.

Characklis, G. W., R. C. Griffin and P. B. Bedient (1999). "Improving the Ability of a Water Market to Efficiently Manage Drought." Water Resources Research **35**(3): 823-832.

Characklis, G. W., B. R. Kirsch, J. Ramsey, K. E. M. Dillard and C. T. Kelley (2006). "Developing Portfolios of Water Supply Transfers." Water Resources Research **42**.

Choi, T. D., O. J. Eslinger, P. Gilmore, A. Patrick, C. T. Kelley and J. M. Gablonsky (1999). IFFCO: Implicit Filtering for Constrained Optimization, Version 2. North Carolina State University, Center for Research in Scientific Computation.

Choi, T. D., O. J. Eslinger, P. Gilmore, A. Patrick, C. T. Kelley and J. M. Gablonsky (1999). IFFCO: Implicit Filtering for Constrained Optimization, Version 2. North Carolina State University, Center for Research in Scientific Computation.

Colby, B. G., K. Crandall and D. B. Bush (1993). "Water Rights Transactions: Market Values and Price Dispersion." Water Resources Research **29**(6): 1565-1572.

DeMouche, L., F. Ward and A. M. Michelsen (2003). Agricultural producers identify barriers to improving water use efficiency in the Rio Grande Basin. Water Security in the 21st Century conference, Univ. Counc. on Water Resour., Washington, D.C.

Easter, K. W., M. W. Rosegrant and A. Dinar, Eds. (1998). Markets for Water: Potential and Performance. Natural Resource Management and Policy. Norwell, MA, Kluwer Academic Publishers.

Efron, B. (1979). "Bootstrap Methods: Another Look at the Jackknife." Annals of Statistics **7**(1): 1-26.

Finkel, D. E. and C. T. Kelley (2004). Convergence analysis of the DIRECT algorithm. N.C. State Univ. at Raleigh.

Fowler, K. R., C. T. Kelley, C. T. Miller, C. E. Kees, R. W. Darwin, J. P. Reese, M. W. Farthing and M. S. C. Reed (2004). "Solution of a well-field design problem with implicit filtering." Optim. Eng. **5**(2): 207-234.

Griffin, R. (1998). The application of water market doctrines in Texas. markets for Water: Potential and Performance. K. W. e. a. Easter. Springer, New York.

Griffin, R. C. and F. O. Boadu (1992). "Water Marketing in Texas: Opportunities for Reform." Natural Resources Journal **32**: 265-288.

Griffin, R. C. and G. W. Characklis (2002). "Issues and Trends in Texas Water Marketing." Water Resources Update **121**: 29-33.

Group, S. C. T. R. W. P. (2000). Water supply options for south central Texas. Austin, Tex. Water Dev. Board.

Guttman, N. B. (1999). "Accepting the standardized precipitation index: a calculation algorithm." Journal of the American Water Resources Association **35**(2): 311-322.

Hamilton, J. R., N. K. Whittlesey and P. Halverson (1989). "Interruptible Water Markets in the Pacific Northwest." American Journal of Agricultural Economics **71**(February): 63-75.

Hamlet, A. F. and D. P. Lettenmaier (1999). "Effects of Climate Change on the Hydrology and Water Resources in the Columbia River Basin." Journal of the American Water Resources Association **35**(6): 1597-1623.

Hearne, R. R. and K. W. Easter (1997). "The economic gains from water markets in Chile." Agricultural Economics **15**: 187-199.

Howe, C. H., D. R. Schurmeier and W. D. Shaw Jr. (1986). "Innovative Approaches to Water Allocation: The Potential for Water Markets." Water Resources Research **22**(4): 439-445.

Howe, C. W. and C. Goemans (2003). "Water Transfers and their Impacts: Lessons from Three Colorado Water Markets." Journal of the American Water Resources Association: 1055-1065.

Howitt, R. E. (1998). Spot Prices, Option Prices, and Water Markets: An Analysis of Emerging Markets in California. Markets for Water: Potentials and Performance. K. W. Easter, Rosengrant, M.W., Dinar A. Norwell. Massachusetts, USA, Kluwer Academic Publishers: 119-140.

Hull, J. C. (1999). Options, Futures, & Other Derivatives. Upper Saddle River, NJ, Prentice-Hall.

Hull, J. C. (1999). Options, Futures, and Other Derivatives. Upper Saddle River, NJ, Prentice-Hall.

IPCC (2007). Working Group II: Impacts, Adaptation, and Vulnerability. Assessment Report 4, Intergovernmental Panel on Climate Change.

Jenkins, M. W. and J. R. Lund (2000). "Integrating Yield and Shortage Management Under Multiple Uncertainties." ASCE Journal of Water Resources Planning and Management **126**(5): 288-297.

Jercich, S. A. (1997). "California's 1995 Water Bank Program: Purchasing Water Supply Options." ASCE Journal of Water Resources Planning and Management **123**(1): 59-65.

- Jettmar, R. U. and G. K. Young (1975). "Hydrologic Estimation and Economic Regret." Water Resources Research **11**(5): 648-656.
- Johnson, H. and D. Shanno (1987). "Option Pricing When the Variance is Changing." Journal of Financial and Quantitative Analysis **22**(2): 143-151.
- Kelley, C. T. (1999). Iterative Methods for Optimization. Philadelphia, SIAM.
- Kirsch, B. R., G. W. Characklis, K. E. M. Dillard and C. T. Kelley (2009). "More Efficient Optimization of Long-term Water Supply Portfolios." Water Resources Research **45**: W03414.
- Kostrowicki, J. and L. Piela (1991). "Diffusion equation method of global minimization: Performance for standard test functions." J. Optim. Theory Appl. **69**(2): 269-284.
- Lall, U. and A. Sharma (1996). "A nearest neighbor bootstrap for resampling hydrologic time series." Water Resources Research **32**(3): 679-693.
- Lavenberg, S. S. and P. D. Welch (1981). "A Perspective on the Use of Control Variables to Increase the Efficiency of Monte Carlo Simulations." Management Science **27**(3): 322-335.
- Lettenmaier, D. P., A. W. Wood, R. N. Palmer, E. F. Wood and E. Z. Stakhiv (1999). "Water Resources Implications of Global Warming: A U.S. Regional Perspective." Climatic Change **43**: 537-573.
- Lund, J. R. and M. Israel (1995). "Optimization of Transfers in Urban Water Supply Planning." Journal of Water Resources Planning and Management **121**(1): 41-48.
- Lund, J. R. and M. Israel (1995). "Optimization of Transfers in Urban Water Supply Planning." ASCE Journal of Water Resources Planning and Management **121**(1): 41-48.
- Lund, J. R. and M. Israel (1995). "Water Transfers in Water Resource Systems." Journal of Water Resources Planning and Management **121**(2): 193-204.
- Matalas, N. C. (1967). "Mathematical assessment of synthetic hydrology." Water Resources Research **3**(4).
- McCarl, B. A., C. R. Dillon, K. O. Keplinger and R. L. Williams (1999). "Limiting Pumping from the Edwards Aquifer: An Economic Investigation of Proposals, Water Markets, and Spring Flow Guarantees." Water Resources Research **35**(4): 1257-1268.
- McKee, T. B., N. J. Doesken and J. Kleist (1993). The relationship of drought frequency and duration to time scales. Eighth Conference on Applied Climatology. Anaheim, California.

McLeod, A. I. and K. W. Hipel (1978). "Preservation of the Rescaled Adjusted Range 1. A Reassessment of the Hurst Phenomenon." Water Resources Research **14**(3): 491-508.

Michelsen, A. M. and R. A. Young (1993). "Optioning Agricultural water Rights for Urban Water Supplies During Drought." American Journal of Agricultural Economics **75**: 1010-1020.

Michelsen, A. M. and R. A. Young (1993). "Optioning Agricultural Water rights for Urban Water Supplies During Drought." American Journal of Agricultural Economics **75**(4): 1010-1020.

More, J. J. and Z. Wu (1997). "Global optimization for distance geometry problems." SIAM J. Optim. **7**(3): 814-836.

Moreau, D. H. and E. E. Pyatt (1970). "Weekly and Monthly Flows in Synthetic Hydrology." Water Resources Research **6**(1): 53-61.

National Research Council, W. S. a. T. B., Division on Earth and Life Studies. (2001). Envisioning the Agenda for Water Resources Research in the Twenty-First Century. Washington, D.C., National Academy Press.

Nieuwoudt, W. L. and R. M. Armitage (2004). "Water Market Transfers in South Africa: Two Case Studies." Water Resources Research **40**.

Palmer, R. and G. W. Characklis (2009). "Reducing the Costs of Meeting Regional Water Demand Through Risk-based Transfer Agreements." Journal of Environmental Management **90**: 1703-1714.

Saliba, B. C. (1987). "Do Water Markets "Work"? Market Transfers and Trade-Offs in the Southern States." Water Resources Research **23**(7): 1113-1122.

Sankarasubramanian, A., U. Lall, F. A. Souza Filho and A. Sharma (2009). "Improved water allocation utilizing probabilistic climate forecasts: Short-term water contracts in a risk management framework." Water Resources Research **45**.

Smith, M. G. and C. M. Marin (1993). "Analysis of Short-run Domestic Water Supply Transfers Under Uncertainty." Water Resources Research **29**(8): 2902-2916.

Srinivas, V. V. and K. Srinivasan (2005). "Hybrid moving block bootstrap for stochastic simulation of multi-site multi-season streamflows." Journal of Hydrology **302**: 307-330.

Stoneking, D. E., G. L. Bilbro, R. J. Trew, P. Gilmore and C. T. Kelley (1992). "Yield optimization using a Gaasprocess simulator coupled to a physical device model." IEEE Trans. Microwave Theory Tech. **40**: 1353-1363.

- Thomas, H. A. and M. B. Fiering (1962). Mathematical synthesis of streamflow sequences for the analysis of river basins by simulation. Design of Water Resource Systems. A. M. e. al. Cambridge, MA, Harvard Univ. Press: 459-493.
- Torczon, V. (1997). "On the convergence of pattern search algorithms." SIAM J. Optim. **7**(1-25).
- VanRheenen, N. T., R. N. Palmer and R. N. Hahn (2003). "Evaluating potential climate change impacts on water resource systems operations: Case studies of Portland, Oregon and the Central Valley, California." Water Resources Update(124): 35-50.
- VanRheenen, N. T., A. W. Wood, R. N. Palmer and D. P. Lettenmaier (2004). "Potential Implications of PCM Climate Change Scenarios for Sacramento - San Joaquin River Basin Hydrology and Water Resources." Climatic Change **62**: 257-281.
- Vaux, H. J. J. and R. E. Howitt (1984). "Managing Water Scarcity: An Evaluation of Interregional Transfers." Water Resources Research **20**(7): 785-792.
- Vicuna, S. and J. A. Dracup (2007). "The evolution of climate change impact studies on hydrology and water resources in California." Climatic Change **82**: 327-350.
- Villinski, M. T. (2004). Valuing Multiple-Exercise Option Contracts: Methodology and Application to Water Markets. American Agricultural Economics Association Annual Meeting, Denver, CO.
- Vogel, R. M. and A. L. Shallcross (1996). "The moving blocks bootstrap versus parametric time series models." Water Resources Research **32**(6): 1875-1882.
- Watermaster's Office (2004). Contract log. W. s. Office. Harlingen, TX, Tex. Commission of Environ. Qual.
- Watkins, D. W. and D. C. McKinney (1999). "Screening Water Supply Options for the Edwards Aquifer Region in Central Texas." ASCE Journal of Water Resources Planning and Management **125**(1): 14-24.
- Watters, P. A. (1995). Efficient Pricing of Water Transfer Options: Non-structural Solutions for Reliable Water Supplies. Economics. Riverside, University of California, Riverside. **Ph. D.**: 145.
- Wilchfort, O. and J. R. Lund (1997). "Shortage Management Modeling for Urban Water Supply Systems." ASCE Journal of Water Resources Planning and Management **123**(4): 250-258.
- Wood, A. W., D. P. Lettenmaier and R. N. Palmer (1997). "Assessing climate change implications for water resources planning." Climatic Change **37**: 203-228.

Young, G. K. and R. U. Jettmar (1976). "Modeling Monthly Hydrologic Persistence." Water Resources Research **12**(5): 829-835.

Young, R. A. (1986). "Why Are There So Few Transactions Among Water Users?" American Journal of Agricultural Economics **68**(6): 1143-1151.