Using Water Transfers to Manage Supply Risk

Gregory W. Characklis[1], Brian R. Kirsch[1], Jocelyn Ramsey[2], Karen E. M. Dillard[3] and C. T. Kelley[3]

Symposium on Safe Drinking Water: Where Science meets Policy

University of North Carolina at Chapel Hill
Chapel Hill, North Carolina

March 2006

[1]Assistant Professor[*] and Research Assistant, respectively, Dept. of Environmental Sciences and Engineering, School of Public Health, Rosenau Hall – CB7431, University of North Carolina, Chapel Hill, NC, 27599-7431, p: (919) 843-5545, f: (919) 966-7911, email: charack@email.unc.edu (G. W. Characklis), bkirsch@email.unc.edu (B. R. Kirsch)
[2]Engineer, Powel Group, 239 Menzies Street, Suite 210, Victoria, BC V8V 2G6, email: jramsey@powelgroup.com
[3]Research Assistant and Professor, respectively, Department of Mathematics, Center for Research in Scientific Computation, Box 8205, North Carolina State University, Raleigh, NC 27695-8205, email: kedillar@unity.ncsu.edu (K. E. M. Dillard), tim_kelley@ncsu.edu (C. T. Kelley)

Using Water Transfers to Manage Supply Risk

ABSTRACT

Most cities currently rely on water supplies with sufficient capacity to meet demand under almost all conditions. However, the rising costs of water supply development make the maintenance of infrequently used excess capacity increasingly expensive, and more utilities are considering the use of water transfers as a means of cost effectively meeting demand. Transfers can take place between utilities, as well as different user groups (e.g., municipal, agricultural), and can involve both treated and untreated water. In cases where both the “buyer” and “seller” draw water from the same supply, contractual agreements alone can facilitate a transfer, but in other cases new infrastructure (e.g., pipelines) will be required. Developing and valuing transfer agreements and/or infrastructure investments will require probabilistic supply/demand analyses that incorporate elements of both hydrology and economics. The complexity of these analyses increases as more sophisticated types of agreements (e. g., options) are considered.

This discussion will revolve around the methods used to develop minimum (expected) cost portfolios of supply assets that meet specified reliability goals. A case study will be described with attention to: the role that transfers can play in reducing average supply costs and tradeoffs between costs and supply reliability. This work involves the development of hydrologic-economic simulations that are linked with an implicit filtering search algorithm designed for the types of “noisy” optimization surfaces that are typically generated when the objective function involves an expected value (cost, in this case). Results will provide insights into the cost savings potential of more flexible water supply strategies.

INTRODUCTION

Rising water demand and concerns over scarcity have driven an increasing number of regions to explore market-based approaches to water resource management. Nonetheless, many water markets remain relatively unsophisticated, with transactions revolving primarily around permanent transfers or multi-year leases While a number of studies have shown that these types of transfers encourage long-term allocation efficiency, such transfers provide a less cost-effective means of managing short-term scarcity. Rising demand 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. This has driven increased interest in the potential efficiency gains associated with "spot" leases and options.

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.

Previous studies have used either linear or stochastic programming techniques to identify combinations of supply alternatives (e.g., infrastructure, transfers and conservation) that minimize the expected costs of meeting urban water demand. 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 supply alternatives. This work expands on earlier studies by employing a simulation-optimization approach that allows for the exploration of some important concerns. Earlier studies, have assumed that a city will acquire water at exactly the time it is needed or, alternatively, that the city has perfect information regarding its future needs when it makes a purchase. Such a scenario is at odds with the behavior 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. Uncertainty with respect to spot market prices is also a concern, therefore spot lease prices are represented as distributions, and this information is used to price options in a risk-neutral manner consistent with financial theory. In addition, while minimizing expected supply costs is important, cost variability will also play a role in decisions regarding a portfolio's suitability, consequently, this work also evaluates tradeoffs between the two.

The modeling approach employed here consists of a hydrologic-market simulation embedded within a search-based optimization algorithm. 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. This approach is applied to the Lower Rio Grande Valley, a region that supports an active water market.

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:

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 .

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. Spot lease prices in each month t are linked to reservoir levels and described as random variables (p_L) (see Figure 1).

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. Option prices (p_O) and exercise prices (p_X) are based on a monthly distribution of spot lease prices (p_L).

Options are priced using a "risk-neutral" approach such that the expected value an option provides relative to a spot lease, does not exceed the option's price. The price (p_O) is calculated by discounting the option's expected value on the call date back to the point at which it was bought, such that

$$p_O = e^{-rT} \bullet E\left[\max\left(\hat{p}_{L_t} - p_X,\ 0\right)\right] \qquad [1]$$

where,

r = discount rate (monthly);

T = 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 that identifies the portfolio that minimizes expected costs while obeying constraints related to reliability and cost variability. The regional context is the western United States, where agricultural water use generally dominates. 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.

Hydrologic-Market Simulation

The simulation runs over a 12 month period, beginning on December 31st ($t = 0$), with the city holding some number of permanent water rights (N_R) 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 (f_{Ro}). In each of the following months, regional hydrologic conditions are simulated using datasets 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 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 generate values for the expected annual cost of the city's portfolio, such that

$$\mathrm{E}[Annual\ Cost] = N_{R_T} p_R + N_O p_O + \mathrm{E}[N_X] p_X + E\left[\sum_{t=0}^{11} N_{L_t} \hat{p}_{L_t}\right] \qquad [2]$$

where,

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

N_O = volume of options purchased to start the year (ac-ft);

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

N_L = volume of spot leases purchased each month (ac-ft).

A water balance is maintained on reservoir storage (R) throughout the simulation such that,

$$R_t = R_{t-1} + i_t - l_t - o_t$$

where,

i_t = volume of reservoir inflows for each month t;

l_t = volume of reservoir losses for each month t;

o_t = volume of reservoir outflows for each month t.

Reservoir inflows available for allocation are calculated as the difference between monthly inflows and losses, multiplied by an instream loss factor (l_I), which accounts for losses incurred between reservoir and user. Inflows are allocated to the city each month (N_{rt}) using a *pro rata* approach based on the total regional number of water rights available (N_R) such that,

$$N_{r_t} = n_t \bullet \left(\frac{N_{R_T}}{\overline{N}_R} \right) \qquad [3]$$

The simulation uses historical records of inflows and city water demand (*d*) to develop a monthly water balance that tracks the city's available supply (*S*). Transfers are assumed to take some time, so the city begins the month with all of the water it will have available in that month, and this amount is compared with a monthly demand value to determine if demand can be met. In each month for which supply does not meet demand, a "*failure*" is recorded, with a special distinction made for critical failures (Supply < 0.7 x Demand) to recognize the more extreme measures a city would need to take under these circumstances. At month's end, information on available supply is combined with knowledge of historical reservoir inflows to estimate future allocations and compute the city's expected supply over the remainder of the year (S_E). Distributions of monthly demand are also used to estimate expected demand over the remainder of the year, and the ratio of these two values is used to make decisions on when and how much water to acquire via leases and/or options (Note: all random variables are identified with a "^").

The decision to acquire water is made by comparing the ratio of expected supply to expected demand against a specified threshold value (α), such that

$$\textit{if} \quad \frac{S_{E_{t+1}}}{\sum_{i=t+1}^{12} \mathrm{E}\left[\hat{d}_i\right]} \leq \alpha \quad \textit{then}, \text{ the city will acquire water,} \quad \text{for } t = 0, 1, 2 \ldots 11 \qquad [4]$$

The question of how much to lease (N_L) and/or exercise (N_X) is made by comparing the ratio of expected supply-to-expected demand with a second specified threshold value (β), such that

$$\frac{\left(N_{L_t} + N_X\right) + S_{E_{t+1}}}{\sum_{i=t+1}^{12} \mathrm{E}\left[\hat{d}_i\right]} = \beta, \quad \text{for } t = 0, 1, 2 \ldots 11 \qquad [5]$$

In all months except the exercise month (t_x), and lease volume purchased is represented as,

$$N_{L_t} = \beta\left(\sum_{i=t+1}^{12} \mathrm{E}\left[\hat{d}_i\right]\right) - S_{E_{t+1}}, \quad \text{for } t \neq t_X. \qquad [6]$$

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). If the lease price is less than the exercise price, the city will simply lease the volume defined above. If, however, the exercise price is less than the lease price, the city will exercise options, with the volume to be exercised expressed as,

$$\textit{if} \;\; \beta\left(\sum_{i=t+1}^{12} \mathrm{E}\left[\hat{d}_i\right]\right) - S_{E_{t+1}} \leq N_O, \;\; \textit{then} \;\; N_X = \beta\left(\sum_{i=t+1}^{12} \mathrm{E}\left[\hat{d}_i\right]\right) - S_{E_{t+1}}, \;\; \textit{otherwise} \;\; N_X = N_O. \qquad [7]$$

Different α and β variables can be specified for individual seasons or even individual months. In the example described, only on parameter pair (α/β) is established for the entire year. 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.

The process of evaluating new allocations and lease/exercise decisions repeats monthly and each annual run within the simulation represents one realization of the expected cost and reliability of a portfolio defined by selected values for the initial conditions (R_0, f_{Ro}) and decision variables (N_R, N_O, α, β). Multiple runs determine a portfolio's expected cost ([2]) and expected reliability, with the latter defined as,

$$E[r_f] = 1 - \left(\frac{\text{failures}}{12 \bullet Years} \right) \qquad [8]$$

where,

r_f = monthly reliability against a failure;

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

A reasonable range of 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, but reliability against these is always maintained > 99.5%.

Simulations can be run over a range of different initial values for permanent rights and options, with surfaces representing the range of portfolio expected costs, as depicted in Figure 2.

Optimization Framework

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

$$\underset{N_R, N_O, \alpha_1, \beta_1, \alpha_2, \beta_2}{\text{Minimize}} \quad Z = E[\mathit{Annual\ Cost}] \qquad [9]$$

Such that:

$E[r_f] \geq$ monthly reliability threshold;

$E[r_{cf}] \geq$ monthly critical reliability threshold.

Figure 2 illustrates a section of the optimization landscape describing expected cost as a function of permanent rights and options. As the volume of leases and exercised options increases the smooth surface becomes more "noisy". 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 increases the computational burden. Implicit Filtering is a finite difference optimization method in which the difference increment (i.e. the size of the finite difference stencil) is varied as the optimization progresses . 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 approach uses a finite difference gradient to compute a search direction for descent, however, unlike classical steepest descent methods in which the negative gradient is used, implicit filtering uses a quasi-Newton model of the Hessian to scale the gradient, a feature that accelerates convergence in the terminal phase of the iteration. In this application, the implementation code, IFFCO (Implicit Filtering For Constrained Optimization) handles constraints in two ways. Simple bound constraints on variables (e.g., $N_0 \geq 0$) 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 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.

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 3). The two reservoirs have a combined storage which is divided between the United States and Mexico according to the treaty of 1944. Since the two reservoirs came on line in 1968, combined U.S. storage in these structures has varied from a low of approximately 0.7 million ac-ft (MAF) to a high of 4.0 MAF. 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. 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 a substantial economic incentive for ag-to-urban water transfers. 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. Lease transactions are relatively simple to complete, and require only that the buyer and seller deliver a one page document to the Watermaster. 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 each year.

RESULTS

The results presented here involve only those in which initial conditions are described as a “worst case” scenario, which includes a January 1 reservoir storage (R_o) is 800,000 ac-ft and an f_{R_o} = 0.2. When permanent transfers are the only means of water transfer (a strategy described as *S0* in Figure 4, with the additional notation that describes it as rights only, *R*), it is assumed that supply cannot be augmented throughout the year. Therefore, the city will need to secure enough water rights to reliably meet its demand under all conditions. For this scenario, the city will need 50,000 ac-ft of permanent rights to meet its water demand with 99.5% reliability (at an annual cost of $1.13 million) and 46,000 ac-ft to meet monthly demand with 98% reliability, at a cost of $1.04 million.

Relative to the situation in which a city must rely on permanent rights to meet demand, short-term transfers would allow a city to pursue more adaptive strategies that can adjust to changing conditions throughout the year. When options are allowed in addition to permanent transfers (strategy *S1*, also noted as “R+O”), the city has more flexibility in managing intra-annual supply and demand variability. A city may hold enough permanent rights to ensure a certain reliability under normal conditions, and acquire options to secure its supply for years in which high water demand and/or drought conditions are experienced. In this type of market, the city can decide how many options to exercise on the exercise date (May 31st), when it has improved information regarding its relative level of supply throughout the rest of the year. Lowering the amount of permanent rights the city owns to 33,400 ac-ft (strategy *S1*) and acquiring 10,000 options reduces the city’s annual supply cost by $0.2 million per year (from $1.13 to $0.93 million), relative to the permanent rights only market (*S0*). If the city’s desired reliability is reduced, then the city’s expected water use declines further, as fewer options are exercised and the number of lapsed options decreases. For example, if the city reduces its desired reliability to 99%, it would need around 5,000 options, and fewer options would be exercised compared to the case with a 99.5% reliability objective. This reduces the water supply cost from \$0.93 to \$0.88 million, a savings of approximately \$56,000 per year

The decision threshold factor, a, has an impact on supply reliability, but its effect on portfolio cost is observed mostly at high levels of reliability (>99%). This factor defines how large the city will allow an expected shortfall to grow before exercising options, therefore, as a increases, the expected shortfall threshold is reduced and reliability increased. When high reliabilities have been reached, increasing this factor no longer continues to affect reliabilities significantly, but still affects cost as it encourages the exercise of more options than are necessary to meet reliability goals

Adding the potential to lease water from the spot market increases the city's flexibility in meeting supply as it may now acquire additional water every month through leasing (as opposed to only exercising options at t_x). Three different strategies are described in this case, each involving the use of rights, options and leases (R+O+L), but with varying levels of permanent rights (Strategies *S1*, *S2* and *S3*). In these strategies, both a and b are significantly lower than in either the "R" or "R+O" markets, indicating that the city can allow the ratio of expected supply-to-expected demand to decline much further before initiating leasing or exercise transactions. In fact, using Strategy S3, in which all water is purchased on the spot market, high levels of reliability can be achieved when both a and b set close to 1.0. Important differences in results can be seen between the strategies analyzed in the rights/options/leases market (R+O+L). With respect to S1, the same number of rights is held (33,400 ac-ft) as in strategy *S1* when only rights and options (R+O) are considered, but the addition of leases leads to a small reduction in expected costs. Strategies *S2* (21,000 ac-ft of permanent rights) and *S3* (0 ac-ft of permanent rights) are less dependent on permanent rights and yield larger expected cost reductions, but cost variability is increased. For example, while meeting 99% reliability using strategy *S0* costs $1.04 million, when all three transfer types are available (i.e. R+O+L) costs decline to $0.87 million for strategy *S1*, $0.73 million for strategy *S2* and $0.55 for strategy *S3*. As would be expected, in the R+O+L market, *S1* requires fewer options than using the same strategy in the R+O market when objective is a 99.5% reliability. However, when the desired reliability is reduced, the number of options bought in the R+O+L market is maintained, and only the number exercised changes. As reliability is reduced, in strategies *S2* and *S3*, the number of options required is either maintained or it declines, and the number of lapsed options varies depending on the amount of leased water (i.e. as more water is leased, less options are exercised). Finally, as the city's dependence on permanent rights is reduced, its expected annual water use (which includes the water that is kept in the city's account and goes unused) is significantly reduced, increasing the overall water availability to the region in wet or average years.

CONCLUSIONS

Most 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. With respect to the solution technique, implicit filtering proves to be an effective search method for the noisy optimization (i.e. expected cost) surface generated in this type of water resource problem, generating repeatable solutions for minimum expected cost and reliability. These results may have important implications for water supply development in the future, particularly in regions where maintaining large volumes of infrequently used supply capacity is either economically or environmentally impractical.

REFERENCES

Black, F., and Scholes, M. (1973). "The Pricing of Options and Corporate Liabilities." *Journal of Political Economy*, 81, 637-659.

Characklis, G. W., Kirsch, B. R., Ramsey, J., Dillard, K. E. M. and C. T. Kelley (2005). “Developing Portfolios of Water Supply Transfers." *Water Resources Research*, (in review).

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

Choi, T. D., and Kelley, C. T. (2000). "Superlinear convergence and implicit filtering." *SIAM Journal of Optimization*, 10, 1149-1162.

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

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., ed., Kluwer Academic Publishers, Norwell. Massachusetts, USA, 119-140.

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

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

Kelley, C. T. (1999). "Iterative Methods for Optimization." *no. 18 in Frontiers in Applied Mathematics*, SIAM, Philadelphia.

Lund, J. R., and Israel, M. (1995). "Optimization of Transfers in Urban Water Supply Planning."

ASCE Journal of Water Resources Planning and Management, 121(1), 41-48.

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

Stoneking, D., Bilbro, G., Trew, R., Gilmore, P., and Kelley, C. T. (1992). "Yield optimization using a GaAsprocess simulator coupled to a physical device model." *IEEE Transactions on Microwave Theory and Techniques*, 40, 1353-63.

Vaux, H. J. J., and Howitt, R. E. (1984). "Managing Water Scarcity: An Evaluation of Interregional Transfers." *Water Resources Research*, 20(7), 785-792.

Watermaster's Office (2004). "Contract Log." Texas Commission of Environmental Quality, Harlingen, Texas.

Watkins, D. W., and McKinney, D. C. (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.

Wilchfort, O., and Lund, J. R. (1997). "Shortage Management Modeling for Urban Water Supply Systems." *ASCE Journal of Water Resources Planning and Management*, 123(4), 250-258.

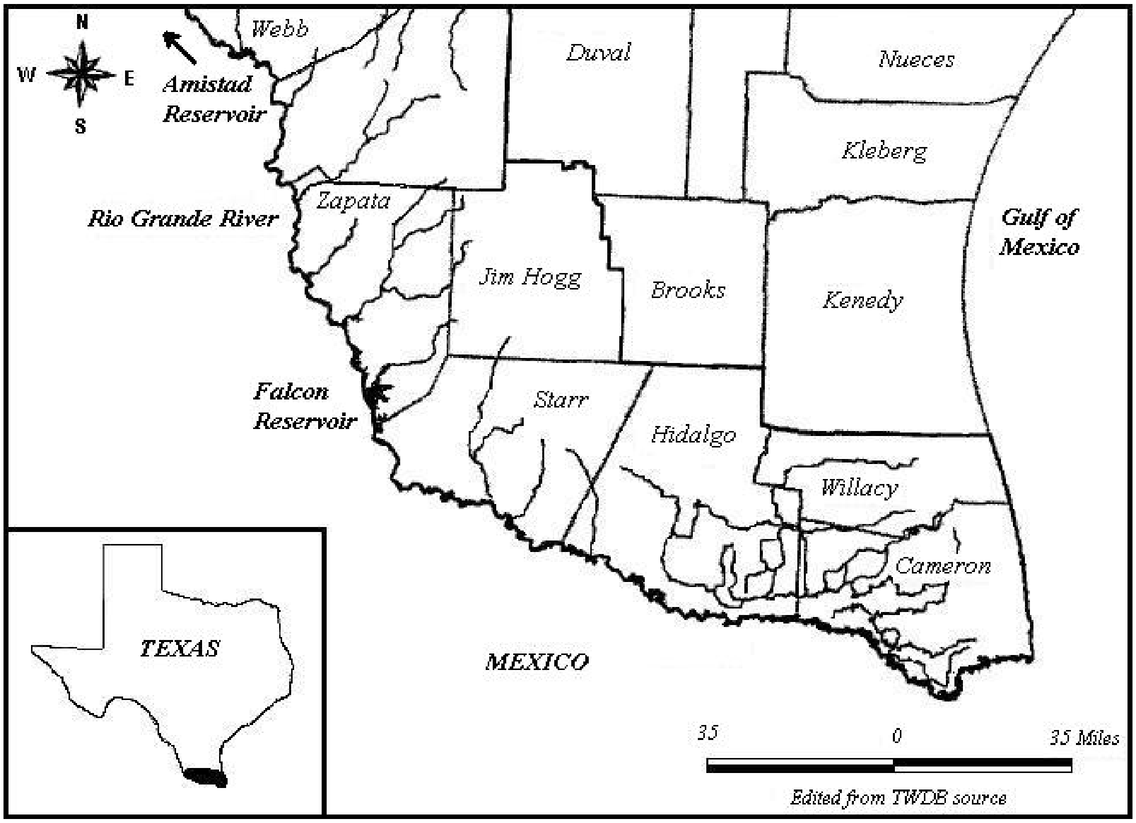

Figure 3   Lower Rio Grande Valley

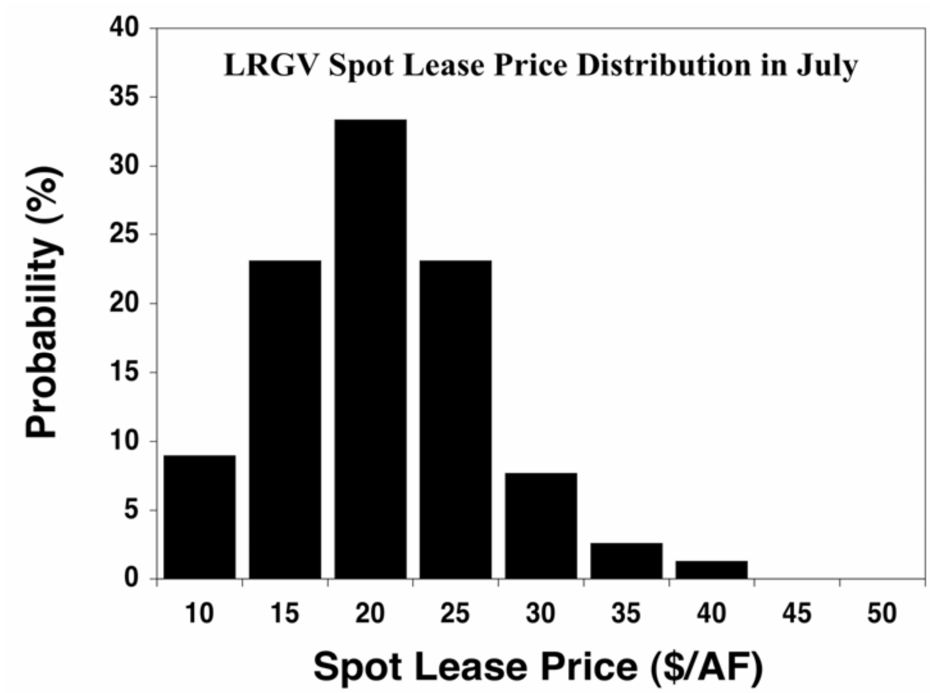

Figure 1   Spot market lease price distribution in the LRGV (1994-2003)

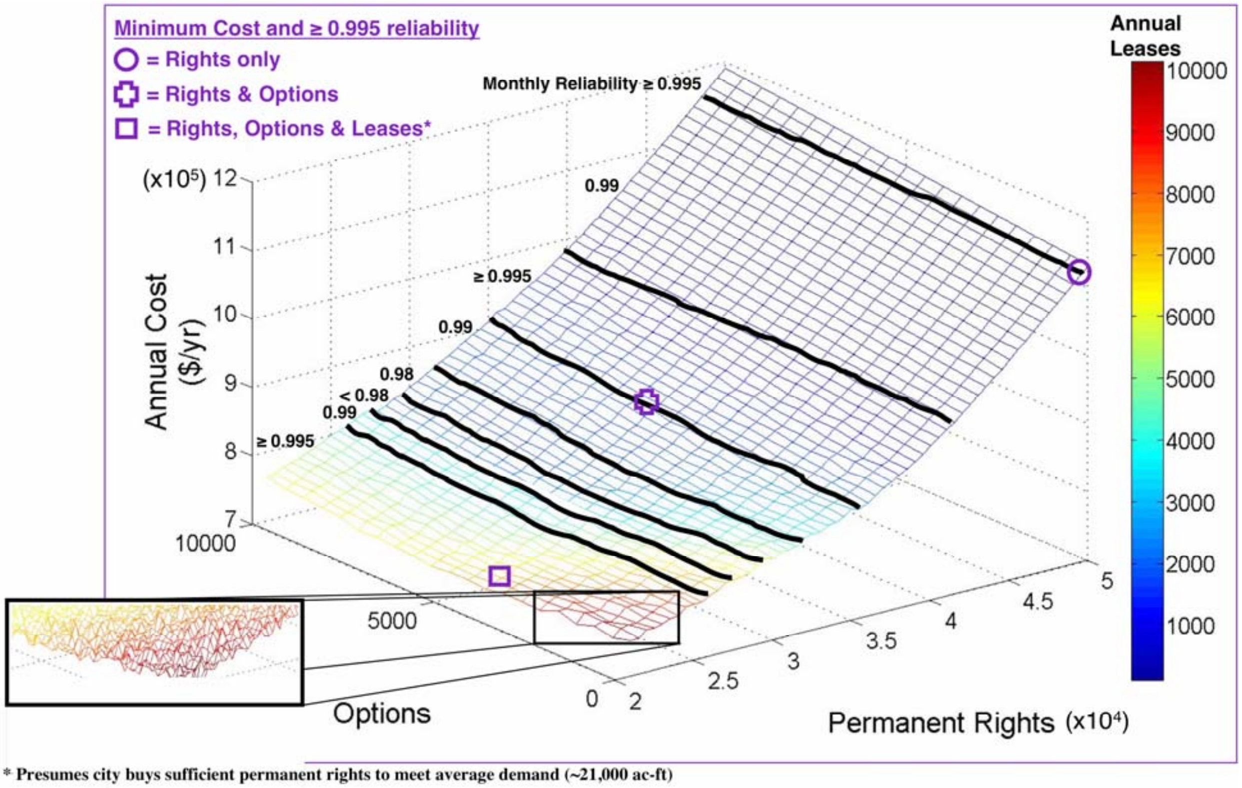

Figure 2 Expected Cost landscape (holding a and b constant)

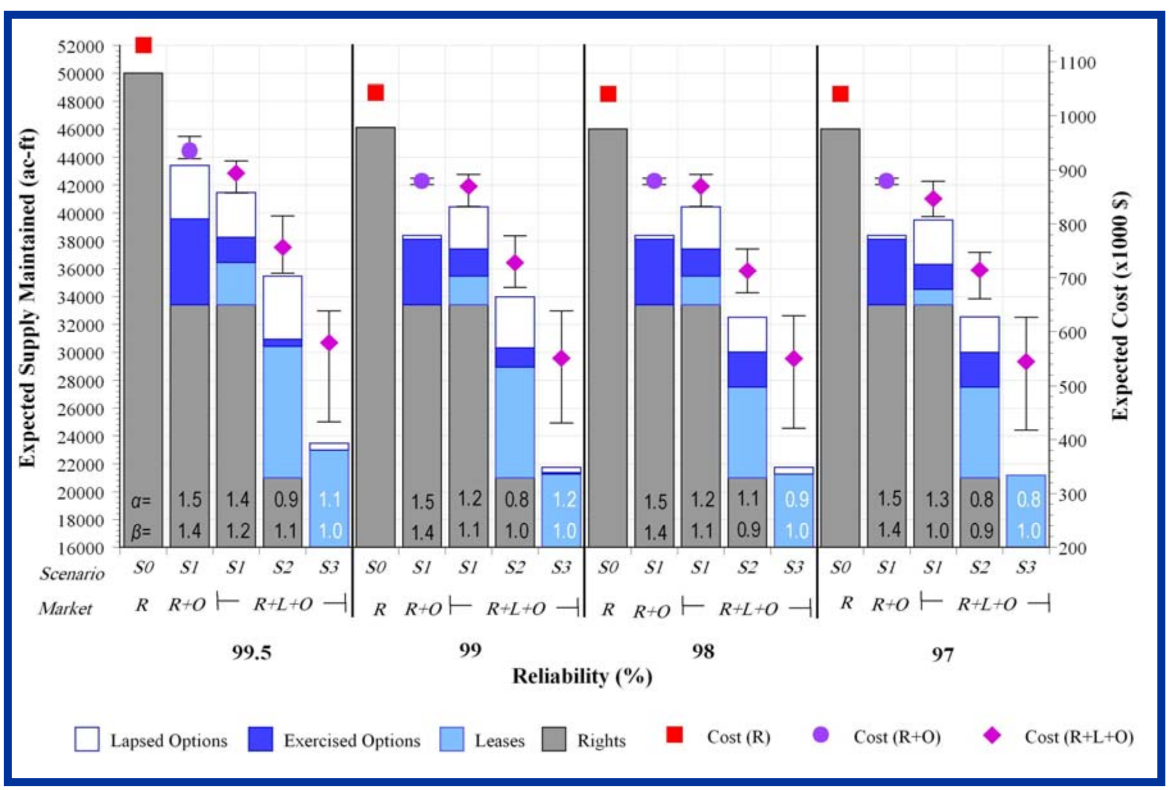

Figure 4 Various Portfolios for a city with an average annual demand of 21,000 ac-ft