Pricing and Inventory Control with Two-Sided Uncertainty: Water Distribution in Southern California*

Claire D. Tomkins† Thomas A. Weber‡

December 2006

PRELIMINARY VERSION

JEL-Classification: D24, D42, D45, D61, Q25, Q58.

Keywords: Allocative Efficiency, Inventory Control, Monopoly Pricing, Water Resources.

*Research in part supported by the Woods Institute for the Environment. We would like to thank Paul Kleindorfer and participants of the 2006 INFORMS Annual Meeting for helpful comments. Any remaining errors are our own.

†Department of Management Science and Engineering, Stanford University. Terman Engineering Center, Stanford University, Stanford, CA 94025-4026. Email: ctomkins@stanford.edu.

‡Department of Management Science and Engineering, Stanford University. Terman Engineering Center, Stanford University, Stanford, CA 94305-4026. Phone: (650) 725-6827. Email: webert@stanford.edu.

Pricing and Inventory Control with Two-Sided Uncertainty: Water Distribution in Southern California

Abstract

We develop an inventory and pricing control model for application to a large intermediary operating in the public sector, the Metropolitan Water District of Southern California (MWD). In doing so, we extend the results from the current literature to the case of two-sided uncertainty with no backlogging, also referred to as "lost sales," and no fixed costs. We show that the optimal policy can be characterized as a base-stock-list-price policy, or (s, S, P)-policy, in keeping with earlier results for the cases of pure demand uncertainty with and without fixed costs and backlogging. The two properties required in our model are, first, independence of the optimal order-up-to-level, y^* or S, and the price, p^* or P, from the inventory state, x (when the constraint $y^* \geq x$ is non-binding) and, second, first-order stochastic dominance of the distribution function of the state transitions, making our exposition considerably more general than those found in earlier papers. We find that, contrary to earlier results, the optimal price can actually be increasing in the inventory state. That is, it is sometimes optimal to increase the price when the constraint $y^* \geq x$ is binding, versus offering a price discount. We treat both profit maximization *and* welfare maximization, where the latter is assumed to be of primary interest to an intermediary in the public sector. Our simulation results reinforce the intuitively appealing conclusion that the optimal order-up-to-level y^* is higher under welfare maximization than under profit maximization, and that the optimal price p^* is lower. The application of the model to MWD provides a number of policy insights, including the differing optimal policy responses to increases in storage, or holding, costs vs. increases in shortage costs, the impact of supply and demand uncertainty on the optimal policy and the payoff function, and the ramifications of correlated demand and supply uncertainty. Our paper aims to both extend the current literature on joint inventory and pricing control models and to draw attention to the application of these models to public pricing problems, an application which has not, to date, received significant attention.

JEL-Classification: D24, D42, D45, D61, Q25, Q58.

Keywords: Allocative Efficiency, Inventory Control, Monopoly Pricing, Water Resources.

1 Introduction

Interest in inventory-control and joint inventory- and pricing control models is longstanding, dating back to pioneering work in the early 1950s by Arrow et al. (1951) and Whitin (1952, 1955). The traditional application of such models has been to the retail sector: a retailer faces the joint decision of how many items to order and what price to set. The components of the model include a unit ordering cost, holding costs associated with storing inventory, shortage costs representative of the cost of filling backlogged demand, and a (possibly uncertain) demand function. The application of these traditional operations and management models to the problem of inventory and pricing control in the public sector has not received significant attention to date. Herein we develop an inventory and pricing control model for application to a large water intermediary in the public sector. In doing so, we dispense with two assumptions common in the current literature. First, we remove the "backlogging assumption," in which all lost demand is backlogged and then filled in the next period. This assumption is unrealistic in the case of water deliveries, where at least demand for municipal and industrial (M&I) water is of an immediate nature and cannot be filled at a later date.[1] The backlogging assumption preserves concavity of the revenue function and, ultimately, of the payoff function. Under such assumptions, Federgruen and Heching (1999) establish that the optimal inventory and pricing policy is a base-stock-list-price policy. That is, there is a base-stock level, b, such that whenever the existing inventory level falls below b, an order is placed to bring the current inventory level up to b. If the current inventory state is greater than b, no order is placed; the list price is that price which is optimal for b.

While helpful in achieving succinct proofs, the concavity assumptions are not critical to the preservation of the form of the optimal policy. The optimality of a base-stock-list-price, or (s, S, P)-policy hinges, rather, on the independence of the control variables from the state of the system – at least, their independence until the constraint that requires the order-up-to-level be no less than the current inventory state kicks in.

Chen et al. (2006) present a "lost sales" model, in which they establish that an (s, S, P)-

[1] The large water intermediary studied herein, the Metropolitan Water District of Southern California, supplies several different classes of water: 'full service treated' water, which is generally for M&I, 'interim agricultural' water, and 'storage replenishment' water. The latter might be more accurately modelled under the assumption that backlogging is permitted, since if a storage facility goes unfilled in this period it will still need replenishment in the next period. Demand for storage water is effectively carried over to the next period. As our focus is on M&I supplies, the no backlogging assumption that we adopt is a more accurate representation of the system. To the extent that M&I water is being pre-ordered and stored, however, the assumption of at least partial backlogging may become tenable.

policy, one in which S denotes the order-up-to level if inventory is below level s, and P is the optimal price associated with S. We, like they, conclude that the form of the optimal policy is a base-stock-list-price, or (s, S, P)-policy. Our approach differs quite significantly from Chen et al. (2006), who base their initial derivations on Petruzzi and Dada (1999). Furthermore, our model does not include fixed ordering costs, as explored by Chen et al. (2006). In contrast to both Chen et al. (2006) and Federgruen and Heching (1999), we impose a limit on the supplier's capacity level and on the maximum order quantity, where, in our model, the latter reflects the size of existing long-term supply contracts.

We also introduce supply uncertainty, in addition to demand uncertainty. The Federgruen and Heching (1999) model considers pure demand uncertainty. Earlier treatments of multi-period pricing and inventory-control models with demand uncertainty include Mills (1959), Zabel (1972), and Thowsen (1975). The addition of supply uncertainty introduces two new parameters – the variance of the supply and the correlation between supply and demand. Under supply uncertainty, we see that, in extreme cases, the optimal order quantity can be driven to zero. The optimal order quantity appears to be (weakly) increasing in the supply uncertainty. We also note that the performance of an adapted policy, i.e., one which is responsive to supply and demand uncertainty vs. just demand uncertainty, is significantly improved.

The base-stock-list-price policy dictates an optimal policy when the constraint that the order-up-to level, or base stock, be less than the current inventory state is *non-binding*. In the case that the constraint is binding, it is optimal not to order any additional items, but the policy does not explicitly dictate the optimal price. Federgruen and Heching (1999) establish that the price is non-increasing in the inventory state, which in turn implies the intuitively appealing policy of price discounting in the case of excess inventory. While we find that in most cases the price is decreasing in the inventory state, we also find, somewhat surprisingly, cases where the price is *increasing* in the inventory state – the latter occurring, for example, in cases where demand is highly inelastic.

Our case study is of the Metropolitan Water District of Southern California (MWD), the largest water intermediary in California, with a current operating budget of over a billion dollars. MWD is a major player in the quasi-monopolistic water distribution in the state. As such, it contracts with the State Water Project (SWP) for the majority of its water (over two-thirds, or 2 million acre feet (maf) annually); contracts for Colorado River water, and various options contracts and fallowing agreements with farmers, as well as conjunctive use agreements, constitute the remainder of MWD's supply. In addition, MWD manages a number of reservoirs and groundwater storage programs.

The intermediary's demand and supply are subject to significant random shocks. On the supply side, the available upstream water resources are largely determined by the precipitation levels and the hydrological conditions in Northern California.[2] In general, both municipal and agricultural water demand increases with temperature.

MWD primarily supplies water to its member agencies. There are currently 26 member agencies – mostly city or county water agencies or authorities, including the Los Angeles Department of Water and Power and the San Diego County Water Authority. The member agencies are voting members of MWD's board, with voting rights proportional to the assessed valuation of the land within the agency's service area. As such, they review the annual rates proposed by MWD's staff. The current rate structure is designed to recuperate average costs.[3] A fixed fee is charged to each member agency, billed as a "Capacity Charge" and a "Readiness-to-Serve Charge," and the remaining costs are recovered through a per-unit charge per af of delivered water.

In our model, we treat both the case of profit maximization and welfare maximization, positing that MWD should be interested in the latter. The former is the standard case in the application of inventory and pricing control models. Welfare maximization follows the tradition of public economics, wherein a "governing board" is assumed to have control over the intermediary's behavior and the power to set prices and regulate quantities (Bös 1985). Public pricing has been widely studied and, as a result, there are a number of available pricing models. One of the most prominent in terms of its application in the public utility sector is the peak-load pricing model, as applied to the electricity sector (Crew, Fernando, and Kleindorfer 1995). The peak-load pricing model dictates optimal prices and capacity, versus order levels: there are no storage, or holding considerations, as the commodity, electricity, is treated as a non-storable commodity.

Our adapted inventory and pricing model with two-sided uncertainty allows for storage, and considers not only the impacts of storage on the intermediary's management, but also the effect of the existing capacity constraints, as short-run capacity is fixed. A limitation of the model is its failure to capture the hydrodynamics of the system. However, as noted above, since the intermediary in our study is not directly diverting water from the source but, rather, contracting for deliveries, this is of less importance. We introduce the possibility of a loss coefficient on supply in our formulation, but do not provide an analysis of its impacts

[2] As MWD is contracting with SWP and not directly diverting water from the source, its supply, as we shall see, is not as closely coupled to the hydrologic cycle and is, rather, affected by the SWP operations.

[3] There is a "Water Rate Stabilization Fund," which helps offset revenue shortages in years of lower deliveries, using the balance from high-demand years.

here.

The water management problem has been studied in detail, both because of its prevalence and importance, with numerous case studies pertaining to California. The bias has been to adopt a systems modelling approach, building models that realistically depict the water network, with flows, capacity constraints, and use values associated with delivery to each "node." See, for example, Draper et al. (2003), Draper and Lund (2004), and Howitt and Vaux (1984).

These models provide critical insight into operations, as well as the value of expanding capacity, reflected in the shadow prices reported for each node. Prices are exogenous in these models, which also usually treat demand as a fixed entity. Our model contributes to the water management problem by examining optimal prices and contract levels *within* an existing operations framework. In doing so, it also expands the existing selection of public pricing models, which has not thus far included inventory control and pricing models from the operations and management literature. The ability of inventory and pricing control models to provide insight into not just the impact of capacity levels, through sensitivity analysis, but also optimal ordering, or contractual, levels, their inclusion of two-sided uncertainty, and their flexibility to treat profit and welfare maximization makes them well-suited to the problem of both public and regulated private utility operations.

Our paper extends the results in the operations and management literature regarding optimal dynamic inventory control policies under uncertainty. Specifically, we extend the case of pure demand uncertainty to that of two-sided (supply and demand) uncertainty, which allows us to also model the correlation between supply and demand. We introduce four other extensions to the model in order to make it more suitable for our intended application: no backlogging of demand ("lost sales"), an explicit capacity level, a loss coefficient on supply, and welfare maximization. The first, no backlogging of demand, or lost sales, has been treated, as discussed by Chen et al. (2006) and, prior to that, by Polatoglu and Sahin (2000). They take a different approach from ours and also include fixed costs in their models. We establish that the form of the optimal policy, as found in previous less general settings, for example, Federgruen and Heching (1999), Polatoglu and Sahin (2000), and Chen et al. (2006), is a base-stock-list-price, or (s, S, P)-policy. We also establish conditions under which, with no underlying concavity assumptions, the optimal price is decreasing *and* increasing in the inventory state, x. In cases of highly elastic demand, for instance, the price may be increasing in x.

We observe, in our study of MWD, that the impact of (positively correlated) dual supply and demand uncertainty is to *decrease* the order-up-to-level and *increase* the list price from that which is optimal under pure demand uncertainty. That is, in a pure demand uncertainty setting and considering the case of very high demand uncertainty (ten times the estimated actual variance), we see that the optimal order-up-to-level reaches the capacity (3.25 maf), with a price between $800 and $850, whereas in the case of pure supply uncertainty and very high supply uncertainty (again, ten times the estimated variance), the optimal order-up-to-level does not reach capacity, staying closer to 2 maf, and the price is above that for the pure demand uncertainty case. Finally, for the case of two-sided uncertainty, the optimal order-up-to-level falls between the case of pure demand and pure supply uncertainty; however, the optimal list price is much higher than in either case. With mixed uncertainty, you charge a much higher price. In general, the optimal order-up-to-level and list price are increasing in demand uncertainty, supply uncertainty, and two-sided uncertainty. We also observe the impacts of the correlation coefficient, concluding that the optimal order-up-to-policy is decreasing in the degree of the correlation, while the list price is unaffected. Finally, we perform sensitivity analysis for the shortage and holding costs in the model, finding that the control levers used by the intermediary to adjust the optimal policy (the order-up-to-quantity vs. the price) differ in the two cases. This underlines the importance of having these dual controls available.

In what follows, we first introduce the model, defining the relevant variables and formulating the maximization problem as a dynamic program. We then proceed to establish a number of properties of the model including the optimal policy and the potential price movement when dependent on the inventory state. Finally, we present the results of the application of the model to the MWD.

2 The Model

Consider a dynamic water distribution system, the state of which is characterized at each time $t \in \{0, 1, \ldots\}$ by the available water inventory $x_t \in X = [0, \bar{x}]$, where $\bar{x} > 0$ corresponds to the total storage capacity. The water distribution system is run by an intermediary, who, at each time t, takes decisions about the conveyance price $p_t \in \mathbb{R}$ and the order-up-to level $y_t \in [x_t, \bar{x}]$. The length of a discrete time interval in the analysis is chosen in accord with the actual schedule of decisions.[4]

[4] In Section 4, we discuss an application where decisions about order contracts and prices are made on an annual basis.

The demand $\tilde{d}_t$ for water at time t is a random variable, the expectation of which is decreasing in the price p_t at time t. For simplicity, we consider here an affine dependence of (nonnegative) demand on price with additive uncertainty, so that

$$\tilde{d}_t(p_t) = [\alpha - \beta p_t + \tilde{\varepsilon}_t]_+ \, ,$$

where α, β are positive real constants and $\tilde{\varepsilon}_0, \tilde{\varepsilon}_1, \ldots \sim N(0, \sigma_\varepsilon^2)$ are i.i.d. zero-mean normal random variables with variance σ_ε^2. The water supply $\tilde{s}_t$ available at time t to satisfy demand is a random variable, the realization of which depends on the order-up-to level y_t,

$$\tilde{s}_t(y_t) = \left[\min \left\{ \lambda y_t + \tilde{\xi}_t, \bar{x} \right\} \right]_+ \, ,$$

where $\lambda \in (0, 1]$ describes a possible transmission loss and $\tilde{\xi}_0, \tilde{\xi}_1, \ldots \sim N(0, \sigma_\xi^2)$ are i.i.d. zero-mean normal random variables with variance σ_ξ^2. The available supply $\tilde{s}_t$ cannot exceed the total storage capacity $\bar{x}$. The quantity of water delivered (or *demand served*) at time t is

$$\tilde{q}_t(p_t, y_t) = \min \left\{ \tilde{d}_t(p_t), \tilde{s}_t(y_t) \right\} .$$

The cost $C(x_t, y_t)$ of delivering water supplies in a given time period t depends on the (nonnegative) difference between order-up-to level y_t and the currently held water inventory x_t:

$$C(x_t, y_t) = c(y_t - x_t)$$

where c denotes a positive unit delivery cost. The unit cost of storing, or holding, water inventory is $h > 0$, while the cost of unmet demand is $r > 0$, so that the expected loss associated with inventory deviations from demand at time t is

$$L(p_t, y_t) = E \left[\max \left\{ h \left(\tilde{s}_t(y_t) - \tilde{d}_t(p_t) \right), r \left(\tilde{d}_t(p_t) - \tilde{s}_t(y_t) \right) \right\} \Big| \, p_t, y_t \right] .$$

In inventory-control models with backlogged demand, the cost of unmet demand, r, is usually measured as the cost of fulfilling backlogged orders. In our model, r has a different interpretation, in keeping with the public pricing and, specifically, peak-load-pricing literature. Here, r is the shortage cost, which is defined in the peak-load-pricing literature as comprised of two parts: the welfare loss associated with curtailed supply and the disruption (or outage) cost, where the latter includes the cost of spoiled goods and lost productivity (Crew, Fernando, and Kleindorfer 1995).[5] The welfare loss is already accounted for in our model, since we are modelling *demand served* rather than demand. Hence, the shortage cost r includes just the outage costs, i.e., loss in productivity and the cost of spoiled goods.

[5] In an electricity outage, the cost of spoiled goods, for example, may include perishable goods, whereas in a period of water shortage it may include perennial crops such as tree crops and/or landscape plants.

We treat both profit maximization and welfare maximization in the model. With regards to the latter, the expected consumer surplus associated with the linear demand function is given as

$$\mathrm{CS}(p_t, y_t) = E[(\hat{p}_t - p_t)\tilde{q}(p_t, y_t) + (\tilde{q}(p_t, y_t))^2/(2\beta)]$$

where $\hat{p}_t(\tilde{q}_t, \tilde{\varepsilon}_t) = (\alpha + \tilde{\varepsilon}_t - \tilde{q}_t)/\beta$ corresponds to a fictitious (random) market clearing price at the (random) quantity $\tilde{q}_t$. This clearing price can never be smaller than the actual price $p_t = (\alpha + \tilde{\varepsilon}_t - \tilde{d}_t)/\beta$, since by definition the demand served cannot exceed the actual demand, whence $\mathrm{CS}(p_t, y_t) \geq E\left[\tilde{q}^2(p_t, y_t) | \, p_t, y_t\right]/(2\beta)$. Let

$$Q_t(p_t, y_t) = E\left[\tilde{q}_t(p_t, y_t) | \, p_t, y_t\right]$$

denote the expected demand served at time t. Given a discount factor $\delta \in (0, 1)$, and a policy $\mu = \{\mu_0, \mu_1, \ldots, \mu_T\}$ with $\mu_t(x_t) = (p_t(x_t), y_t(x_t))$ the intermediary's total discounted expected payoff at time $t \in \{0, \ldots, T-1\}$ is recursively defined as follows:

$$\Pi_T(x_T; \mu) = p_T Q_T(p_T, y_T) - c(y_T - x_T) - L(p_T, y_T) + \omega \, \mathrm{CS}(p_T, y_T),$$

for $t = T$, where

$$\omega = \begin{cases} 1, & \text{if the intermediary is welfare-maximizing,} \\ 0, & \text{otherwise,} \end{cases}$$

and

$$\Pi_t(x_t; \mu) = \left[p_t Q_t(p_t, y_t) - c(y_t - x_t) - L(p_t, y_t) + \omega \, \mathrm{CS}(p_t, y_t) + \delta E\left[\Pi_{t+1}(\tilde{x}_{t+1}; \mu) | \, p_t, y_t\right]\right]\Big|_{(p_t, y_t) = \mu(x_t)}$$

for $t \in \{0, \ldots, T-1\}$, where the stochastic state transition from any x_t to a (random) state $\tilde{x}_{t+1}$ is defined by

$$\tilde{x}_{t+1} = \max\left\{0, \tilde{s}_t(y_t) - \tilde{d}_t(p_t)\right\}.$$

Optimizing over all admissible policies, the associated discrete-time dynamic-programming equation is

$$\Pi_t(x_t) = c x_t + \max_{(p_t, y_t) \in \mathbb{R} \times [x_t, \bar{x}]} \left\{p_t Q_t(p_t, y_t) - c y_t - L(p_t, y_t) + \omega \mathrm{CS}(p_t, y_t) + \delta E\left[\Pi_{t+1}(\tilde{x}_{t+1}) | \, p_t, y_t\right]\right\},$$

for all $t \in \{0, \ldots, T-1\}$, and

$$\Pi_T(x_T) = c x_T + \max_{(p_T, y_T) \in \mathbb{R} \times [x_T, \bar{x}]} \left\{p_T Q_T(p_T, y_T) - c y_T - L(p_T, y_T) + \omega \mathrm{CS}(p_t, y_t)\right\}$$

for $t = T$.

The distribution of transition probabilities from any inventory state x_t to the inventory state $x_{t+1} = x \in [0, \bar{x})$, as a function of choice variables (p_t, y_t), is

$$F(x|p_t, y_t) = P(\tilde{x}_{t+1} \leq x|p_t, y_t) = \int_{-\infty}^{\infty} \left(\int_{\min\{\lambda y_t + \xi, \bar{x}\} - (\alpha - \beta p_t) - x}^{\infty} \varphi(\varepsilon, \xi) d\varepsilon \right) d\xi,$$

where φ denotes the joint density for the possibly correlated demand and supply uncertainties. For concreteness, we consider in this model a bivariate normal distribution with standard deviations $\sigma_\varepsilon, \sigma_\xi > 0$ and correlation $\rho \in (-1, 1)$,

$$\varphi(\varepsilon, \xi) = \frac{\exp\left[-\frac{1}{2(1-\rho^2)} \left(\frac{\varepsilon^2}{\sigma_\varepsilon^2} - \frac{2\rho\varepsilon\xi}{\sigma_\varepsilon \sigma_\xi} + \frac{\xi^2}{\sigma_\xi^2} \right) \right]}{2\pi\sigma_\varepsilon\sigma_\xi\sqrt{1-\rho^2}}.$$

For $x = \bar{x}$ we obtain $F(\bar{x}|p_t, y_t) = 1$, so that F may be discontinuous at the upper bound. Correspondingly, we find for the probability density for the state x_{t+1} that

$$f(x|p_t, y_t) = \int_{-\infty}^{\infty} \varphi\left(\min\{\lambda y_t + \xi, \bar{x}\} - (\alpha - \beta p_t) - x, \xi\right) d\xi$$

for all $x \in (0, \bar{x})$ with atoms at the interval ends of size $F(0|p_t, y_t)$ and $1 - F(\bar{x}|p_t, y_t)$ respectively. We can now rewrite the expected next-period payoff in the above dynamic-programming equation in the form

$$E\left[\Pi_{t+1}(\tilde{x}_{t+1})|\, p_t, y_t\right] = \int_0^{\bar{x}} \Pi_{t+1}(x) f(x|p_t, y_t) dx + F(0|p_t, y_t)\Pi_{t+1}(0) + (1 - F(\bar{x}|p_t, y_t))\Pi_{t+1}(\bar{x}),$$

for all $t \in \{0, \ldots, T-1\}$.

Remark 1 When discretizing the state space to $N+1$ inventory levels $\hat{x}^k = k\bar{x}/N$ for $k \in \{0, \ldots, N\}$, the corresponding transition probabilities from $x_t = \hat{x}^j$ to $x_{t+1} = \hat{x}^k$ are given approximately by

$$\pi_{j,k}(p_t, y_t) = F((2k+1)\bar{x}/(2N)|p_t, y_t) - F((2k-1)\bar{x}/(2N)|p_t, y_t)$$

for $k \in \{1, \ldots, N-1\}$, $\pi_{j,N}(p_t, y_t) = 1 - F((2N-1)\bar{x}/(2N)|p_t, y_t)$, and $\pi_{j,0}(p_t, y_t) = F(\bar{x}/(2N)|p_t, y_t)$. Note that all the transition probabilities are independent of the state x_t at time t, i.e., independent of j, which simplifies computations substantially. In the discretized state-variables, the expected value in the dynamic-programming equation becomes (for all $t \in \{0, \ldots, T-1\}$ and all $k \in \{0, \ldots, N\}$)

$$E\left[\Pi_{t+1}(\tilde{x}_{t+1})|\, p_t, y_t\right] = \sum_{k=0}^{N} \pi_k(p_t, y_t)\Pi_{t+1}(k\bar{x}/N),$$

where, for simplicity, we have written π_k instead of $\pi_{j,k}$. □

3 Optimal Inventory and Pricing Control Policy

We now characterize the optimal dynamic inventory-control and pricing policy. For this it is useful to start by considering the comparative statics of the intermediary's expected per-period payoff

$$g(p, y) + cx = pQ(p, y) - L(p, y) - c(y - x) + \omega \text{CS}(p, y),$$

where, for convenience, we drop the dependence of variables on the time t unless there can be misunderstandings.

Proposition 1 *The expected demand served $Q(p, y)$ and the expected consumer surplus $\text{CS}(p, y)$ are both decreasing in p, increasing in y, and submodular in (p, y).*

Since water is a normal good, a higher price is unambiguously bad news for consumers, so that both expected demand served and consumer surplus are necessarily decreasing in price. The proof of Proposition 1 shows that the marginal decrease of expected demand served Q with respect to price is proportional to the demand-elasticity parameter β and the probability that supply exceeds demand,

$$Q_p(p, y) = -\beta P(\tilde{d}_t(p) \leq \tilde{s}_t(y)) < 0.$$

On the other hand, the marginal increase of Q with respect to the order-up-to level y is proportional to the loss factor λ and the probability that supply neither reaches capacity nor fully covers demand,

$$Q_y(p, y) = \lambda P(\tilde{s}_t(y) < \min\{\tilde{d}_t(p), \bar{x}\}) > 0.$$

Moreover, the submodularity of Q implies that Q_p decreases with an increase in y and, vice versa, Q_y decreases with an increase in p. The marginal decrease of expected consumer surplus CS with respect to a unit price increase is proportional to the expected demand served,

$$\text{CS}_p(p, y) = -Q(p, y) < 0.$$

The marginal increase of CS with respect to a unit increase in the order-up-to level y is proportional to the expected excess demand, conditional on the fact that supply neither reaches capacity nor fully covers demand,

$$\text{CS}_y(p, y) = \frac{\lambda}{\beta} E\left[\tilde{d} - \tilde{s} \,\middle|\, \tilde{s} < \min\{\tilde{d}, \bar{x}\}\right] P(\tilde{s} < \min\{\tilde{d}, \bar{x}\}) > 0.$$

The nonnegative expected losses $L(p, y)$ are not monotonic with respect to either price or order-up-to level. For any given positive holding and shortage cost parameters h and r, they vanish if and only if demand is exactly met by supply. The marginal losses with respect to a unit change in price can be determined analogously to the computations in the proof of Proposition 1,

$$L_p(p, y) = -\beta r + \beta(r + h)P(\tilde{d}_t(p) \leq \tilde{s}_t(y))$$

Similarly,

$$L_y(p, y) = -\lambda r + \lambda(h + r)P(\tilde{d}_t(p) \leq \tilde{s}_t(y))$$

and

$$L_{py}(p, y) = \beta\lambda(r + h) \int_{-\infty}^{\bar{x}-\lambda y} \varphi(\lambda y + \xi - (\alpha - \beta p), \xi) d\xi > 0,$$

so that the term $-L(p, y)$ in the intermediary's per-period payoff is submodular with respect to (p, y).

At any time t both, a higher price and a higher order-up-to level tends to increase the stochastic next-period inventory level. The decision vector (p_t, y_t) induces a first-order stochastic dominance (FOSD) order of the distributions of $\tilde{x}_{t+1}$.

Proposition 2 *The distributions $F(\cdot|p, y)$ corresponding to the state-transition probabilities conditional on (p, q) are FOSD-ordered in the sense that*

$$(p, y) \leq (\hat{p}, \hat{y}) \quad \Rightarrow \quad F(\cdot|p, y) \preceq_{\text{FOSD}} F(\cdot|\hat{p}, \hat{y})$$

for all $(p, y), (\hat{p}, \hat{y}) \in \mathbb{R} \times X$.

The probability of transiting from a given inventory state to any other inventory state only depends on the chosen order-up-to level y_t and the price p_t. In particular, it is independent of the current inventory state x_t, so that the expected next-period payoff $E[\Pi_{t+1}(\tilde{x}_{t+1})|p_t, y_t]$ is independent of x_t. Since an increase in x_t further constraints the objective, the function $\Pi_t(x) - cx$ must be nonincreasing in x.

Proposition 3 *For any $t \in \{0, \ldots, T\}$ the payoff function $\Pi_t(x)$ is such that*

$$x, \hat{x} \in X \text{ with } x \leq \hat{x} \quad \Rightarrow \quad \Pi_t(\hat{x}) - \Pi_t(x) \leq c\,(\hat{x} - x)\,.$$

The last two results imply the following monotonicity property of the optimal expected payoffs.

Proposition 4 *For any $t \in \{0, \ldots, T\}$ the expected payoff $\bar{\Pi}_t(p, y) = E\left[\Pi_t(\tilde{x}_{t+1}) | \, p, y\right]$ satisfies*

$$(p, y) \leq (\hat{p}, \hat{y}) \quad \Rightarrow \quad \bar{\Pi}_t(\hat{p}, \hat{y}) - \bar{\Pi}_t(p, y) \leq c\left(E\left[\tilde{x}_{t+1} | \, \hat{p}, \hat{y}\right] - E\left[\tilde{x}_{t+1} | \, p, y\right]\right) \leq 0 \qquad (1)$$

for all $(p, y), (\hat{p}, \hat{y}) \in \mathbb{R} \times X$.

We now establish the existence of a nonzero, finite optimal price. The idea is to show that for any fixed order-up-to level y the intermediary's payoffs are 'coercive' in price, in the sense that it is always better for the intermediary to charge a price larger than zero and less than infinity, implying the existence of an optimal finite price.

Remark 2 The intermediary's expected gross revenues $pQ(p, y)$ are generally not concave in (p, y), even though the expected demand served is. To see the latter, note that

$$\begin{aligned} Q_{pp}(p, y) &= -\beta^2 \int_{\mathbb{R}} \varphi(\min\{\lambda y + \xi, \bar{x}\} - (\alpha - \beta p), \xi) d\xi < 0, \\ Q_{py}(p, y) &= -\beta\lambda \int_{-\infty}^{\bar{x} - \lambda y} \varphi(\lambda y + \xi - (\alpha - \beta p), \xi) d\xi < 0, \\ Q_{yy}(p, y) &= -\lambda^2 \int_{-\infty}^{\bar{x} - \lambda y} \varphi(\lambda y + \xi - (\alpha - \beta p), \xi) d\xi < 0. \end{aligned}$$

Let $H(Q) = \partial^2 Q(p, y)/\partial(p, y)^2$ be the Hessian matrix for Q. Then $\det H(Q) = Q_{pp}Q_{yy} - Q_{py}^2 < 0$ and $\operatorname{trace} H(Q) = Q_{pp} + Q_{yy} < 0$, so that both eigenvalues of $H(Q)$ must be negative. This implies that $H(Q)$ is negative definite and thus Q strictly concave in (p, y). However, it can be shown using an analogous method that the function $pQ(p, y)$ may not be concave in (p, y). □

Proposition 5 *For any $t \in \{0, \ldots, T\}$ and any $x \in [0, \bar{x}]$ the optimal price $p_t^*(x)$ lies in $[0, \bar{p}(x)]$, where $\hat{p}(x) = \inf\left\{p \geq 0 : P(\tilde{d}_t(p) \leq \tilde{s}_t(x)) \geq 1/2\right\}$ and $\bar{p}(x) = \max\{\hat{p}(x), r + h - \omega\left(\frac{\alpha}{\beta} - \frac{2\zeta}{\beta}\right) + \rho + \frac{2\bar{x}}{\beta}\}$, where*

$$\zeta = \left(E[\tilde{\varepsilon} | \tilde{d}_t(p) < \tilde{s}_t(y)] - E[\tilde{q} | \tilde{d}_t(p) < \tilde{s}_t(y)]\right) P(\tilde{d}_t(p) < \tilde{s}_t(y)).$$

The upper bound $\bar{p}(x)$ is nonincreasing in x.

The last result provides an explicit upper bound for the optimal price, corresponding to a price level at which the reduction in demand outweighs all potential benefits of a profit-maximizing intermediary. It is clear that because of $\text{CS}_p < 0$ a welfare-maximizing intermediary never charges a higher price than a profit-maximizing intermediary. As the available inventory x increases, the upper bound $\bar{p}(x)$ for the price weakly decreases.

Proposition 6 *(i) Under welfare maximization* ($\omega = 1$)*, the intermediary's per-period payoff is submodular in* (p, y)*. (ii) Under profit maximization* ($\omega = 0$)*, the intermediary's per-period payoff may become supermodular in* (p, y) *provided that demand is sufficiently inelastic, i.e., if* β *is small.*

The submodularity and supermodularity results come from examination of $g_{py}(p, y)$, where in the case of welfare maximization, $g_{py}(p, y)$ is seen to be always negative. The negative impact of an increase in price - which decreases expected demand served and, hence, decreases expected consumer surplus as well as expected revenue - and the cost to the intermediary of ordering and storing an additional unit of inventory outweighs the expected gain in revenue and consumer surplus from an increase in y. However, under profit maximization when the possibility that the demand is (largely) unaffected by a price increase exists, i.e., for $\beta << 1$, $g_{py}(p, y) > 0$. The sub- and supermodularity results imply the following Proposition.

Proposition 7 *Let* $t \in \{0, \ldots, T\}$*. (i) Under welfare maximization* ($\omega = 1$)*, the optimal price* $p_t^*(x)$ *is nonincreasing in* x*. (ii) Under profit maximization* ($\omega = 0$)*, the optimal price* $p_t^*(x)$ *is generically nonmonotonic; in particular it can be locally increasing in* x*.*

Part (ii) of the last result is somewhat counterintuitive, as it seemingly contradicts extant results in the literature (see e.g., Federgruen and Heching (1999)). Consider the last period $t = T$ in which the intermediary cares only about the per-period payoff $g(p, x) + cx$, provided that $y^*(x) = x$. When demand is extremely inelastic (i.e., for small β), then the positive marginal benefit of a price increase for the intermediary's payoff increases with the inventory level, since that price increase can be paid by a larger number of consumers. Figure 3 shows a numerical example, in which – to obtain the effect at a reasonable price – the optimal price was held close to a target level p_{target} by adding a negative quadratic term $-k(p - p_{\text{target}})^2$, for some $k > 0$, to the intermediary's per-period payoffs, which penalizes deviations from the target price, and does not influence the supermodularity properties of g.

Proposition 8 *For any* $t \in \{0, \ldots, T\}$ *the optimal order-up-to policy* $y_t^* : X \to X$ *is such that*

$$y_t^*(x) \in \{x, \bar{y}_t^*\}$$

for all $x \in X$*. The optimal pricing policy is a list-price-policy as long as* $y_t^* = \bar{y}_t^*$*.*

Proposition 8 shows that the results in the operations management literature on optimal inventory-control and pricing policies under uncertainty continue to hold in the case of two-sided uncertainty with lost sales (i.e., without demand backlogging): it is optimal for the

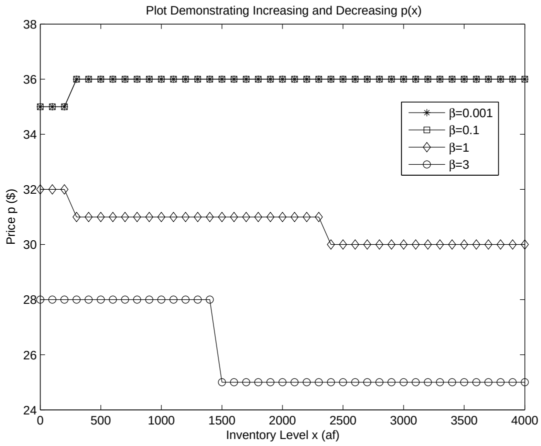

Figure 1: Example of Increasing $p^*(x)$, with $p_{\text{target}} = 20$, $k = 6$. The other parameters are $\alpha = 174$, $\beta = 3$, $c = 22$, $h = 0.20$, $r = 22$, and $\bar{x} = 400$. For $\beta \leq 0.1$, the optimal price $p^*(x)$ increases from \$35 to \$36 as the inventory level x increases from 0 to 500 af.

intermediary to follow a base-stock-list-price policy, also termed (s, S, P)-policy, in which the intermediary orders up to a base-stock level y^* whenever the current water inventory x is below y^*, and does not order otherwise, and the price p^* is the optimal price when y^* is the current inventory state *and* the constraint that $y^* \geq x$ is non-binding. Note that this is a particular form of (s, S, P), where $s = S = y^*$. That is, the intermediary orders up to y^* whenever the inventory level x is below y^* setting a price $P = p^*$, and otherwise does not order any water and uses a price $p^*(x)$ that generally depends on x. The latter dependence is, at least under profit maximization, due to Proposition 6 (ii) generically nonmonotonic.

This characterization hinges on two properties: first, in the case where the constraint that the order-up-to-level cannot be less than the inventory state, i.e., $y^* \geq x$, is *non-binding*, the solution hinges on the independence of the optimal order-up-to-level and price, (p^*, y^*), from the inventory state, x and, second, in the case where the constraint is *binding*, the solution depends on the expected next-period payoff being non-increasing in y, which ensures that $y_t^*(x) \in \{x, \bar{y}_t^*\}$. The first property, the independence of (p^*, y^*) from the state, x, can be seen through examination of the derivatives of $g(p, y)$. The second property is established in Proposition 4 above.

These properties are satisfied under general conditions. The FOSD of the distribution function for the state transitions, $F(\cdot|p, y) \preceq_{\text{FOSD}} F(\cdot|\hat{p}, \hat{y})$, and Proposition 3 are needed to establish Proposition 4. Proposition 3 requires only the independence of the optimal policy from the inventory state, x. Hence, we see the FOSD of $F(\cdot|p, y)$ and the independence of the optimal policy from the state of the system (when the constraint $y^* \geq x$ is non-binding) are the general conditions, for which the characterization holds true. Concavity assumptions and restrictions on the forms of the demand and loss functions, beyond ensuring independence from the state x, are not integral to the characterization of the optimal policy.

4 Application of the Model to the Metropolitan Water District of Southern California

As mentioned in Section 1, the Metropolitan Water District of Southern California contracts with the State Water Project for close to 2 maf of water a year. Although MWD holds other supply contracts, the SWP contract is the largest and currently accounts for over two-thirds of MWD's contractual supply. In dry years or when reservoir levels are low, the SWP deliveries are cut by a percentage, as determined by the state's Department of Water Resources (DWR). Figure 2 (b) shows MWD's total estimated supply over a 20-year period, from 1986 to 2005. Note that we have assumed that the other contractual supplies have instituted the same cuts in supply as DWR; that is, if DWR cuts SWP supply by 25%, then it is assumed that the remaining one-third of MWD's contractual supply from other sources was also cut by 25%.[6]

Demand for water in Southern California is higher during dry years; our regression analysis indicated that the uptake in demand is present not just for years categorized as "critical" or "dry," but also for those categorized as "below normal" in the DWR's five-point year-type index. (The other two year classifications in the index are "above normal" and "wet.") We would therefore expect to find a negative correlation coefficient supply and demand. However, the DWR does not simply divert water from the major river systems in northern California to provide to its contractors, such as SWP. Rather, it operates a system of reservoirs, which it uses to offset supply shortfalls in dry years. Recognizing that demand by all sectors will increase, the SWP will not necessarily cut supplies in a single dry year, but instead opt to instead release water from its reservoirs; on the other hand, when facing a multi-year dry spell, DWR may need to cut supplies in order to maintaining minimum reservoir levels. As a consequence, Figure 3 (a) shows that the correlation coefficient between MWD demand and supply is positive. Figure 3 (b) on the other hand shows that the "true" correlation

[6]Under this assumption, the supply variability is overestimated. In actuality, MWD's additional contractual supply is not perfectly correlated with the SWP supply and, hence, at times where SWP imposes cuts in deliveries to MWD, the remaining contracts may still be filled.

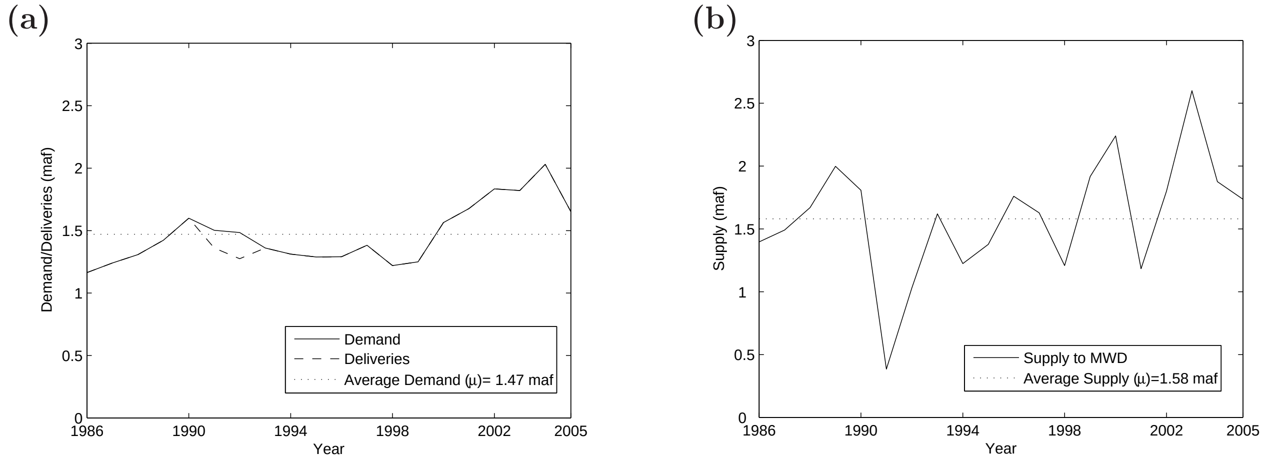

20-Year Demand 20-Year Supply

Figure 2: Average supply slightly exceeds average demand; however, supply is also more variable.

between supply and demand is indeed negative.

The early 1990s was the culmination of a six-year drought in California, the worst drought on record in a 100-year period. In 1991 and 1992, supply levels fell drastically, and MWD implemented rationing. Rationing was instituted in accordance with MWD's five-stage Incremental Interruption and Conservation Plan (IICP), where Stage 3 required that all member agencies reduce their stated demands to 10% below 1990 levels; Stage 4 required a 20% reduction from 1990 levels, and Stage 5, declared in April 1991, 30%.[7] (See Appendix B for further details of the IICP.) The line "Demand" in Figure 2 (a) estimates actual demand vs. deliveries, during the period when rationing was in place (from 1991 to 1993), where all demands by individual member agencies have been adjusted back to the 1990 levels whenever the constraint that demand be less than 10%, 20% or 30% of 1990 demand levels, depending on the level of rationing in place for that month, are binding for the agency.

The estimates for α and β, the market potential and the elasticity-of-demand parameter, respectively, in our linear demand model were estimated using a multiple linear regression, with $R^2 = 0.8652$. The model accounted for seasonality and population growth in Southern California, using the population growth in San Diego county, as reported in the U.S. census,

[7]MWD, letters to the Board containing the subject, "Incremental Interruption and Conservation Plan," Nov. 13, 1990 to Sept. 9, 1993, `www.mwdh2o.com`, currently available at `http://edmsidm.mwdh2o.com/idmweb/search.asp`, with "Subject contains:" set equal to "Incremental Interruption and Conservation Plan."

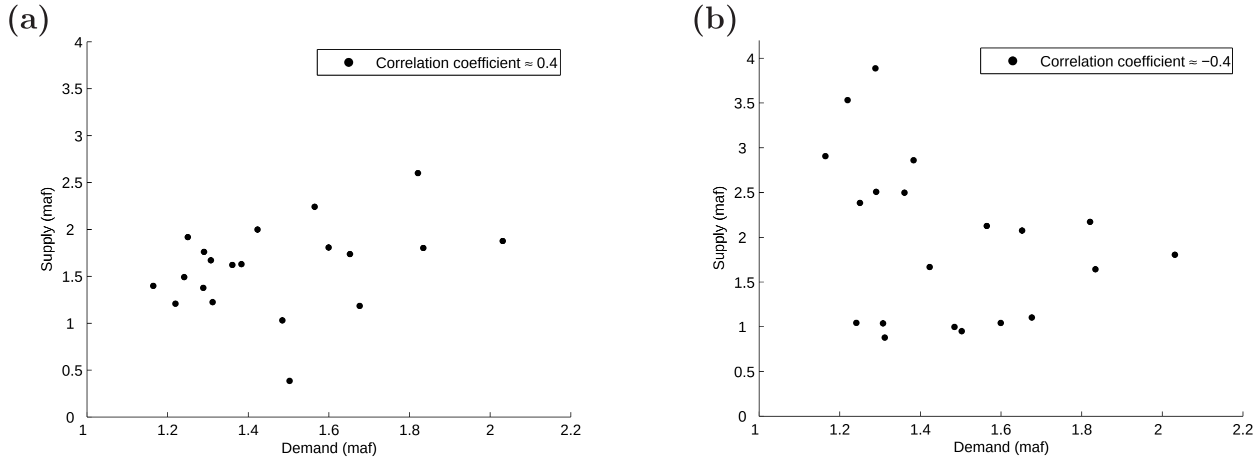

Supply from SWP vs. Demand

Supply from Sac R. vs. Demand

Figure 3: Actual supply to MWD, from SWP, is positively correlated with demand, whereas supply from the source (the Sacramento River) is negatively correlated with demand.

as the proxy for regional population growth. Not all areas in the MWD service area are growing as rapidly as San Diego county; however, there is a strong regional growth trend. (San Diego is currently the largest user of MWD water, accounting for roughly 30% of their total deliveries in 2005.) The seasonality factors are accounted for by the dummy variables dcrit, which is equal to one if the year-type is "dry," "critical," or "above normal," and zero otherwise, and dwet, which is equal to one if the year type is "wet," and zero otherwise. The variable cumpopgrowth tracks the cumulative population growth each year in San Diego County.[8] The estimated values on the coefficients are reported in Table 1.

For the simulation runs, population growth was not included in the model, and, hence the dummy variable cumpopgrowth below was assumed equal to zero in all scenarios; similarly, the dummy variables dcrit and dwet were set to zero. Under these assumptions, the demand model is $d = \alpha' - \beta p + \epsilon$, as in our derivations, where α' is the model-estimated α plus the most recent population factor, which is held constant in our simulations.

The unit cost of supplying water, c, is estimated using MWD's records for water supply in 2005, as their total annual costs, minus the fixed costs (as charged through their Capacity Charge and Readiness-to-Serve), divided by the total water delivered.[9] The holding cost

[8] U.S. Census Bureau (1986–2005) "Annual Estimates of the Population for Counties," www.census.gov.

[9] MWD, letter to the Board, 10 Jan. 2006, Board Letter Archives, www.mwdh2o.com, currently at http://edmsidm.mwdh2o.com/idmweb/cache/MWD%20EDMS/003678343-1.pdf.

α	β	γ_c	γ_w	g
1,749,000 af (160,890)	2,600 af/$ (740)	17,900 (57,620)	-122,100 (66,350)	4.4 (0.77)

Table 1: Parameters for the demand model $d(p) = \alpha - \beta p + dcrit\ \gamma_c + dwet\ \gamma_w + cumpopgrowth\ g$ with standard errors (in parentheses) reported for the 5% significance level.

for water, which is negligible for an additional af of water in a reservoir but is estimated to be $140/af for groundwater storage, is the weighted average of these two, where the weights are the percentage of existing reservoir capacity (64%) and groundwater capacity (46%), respectively.

Description	**Symbol**	**Value**
Market potential	α	2.9392 maf
Elasticity of demand	β	2.6 taf
Unit order cost	c	$400
Holding (storage) cost	h	$0 to $64/af
Shortage cost	r	$0 to $5,885/af
Capacity	$\bar{x}$	3.26 maf
Loss coefficient	λ	1
Discount factor	δ	0.95
Joint dist. of supply and demand	$\varphi(\varepsilon, \xi)$	Bivariate normal
Std. dev. of demand	σ_ϵ	2.372 taf
Std. dev. of supply	σ_ξ	4.632 to 8.895 taf
Correlation coefficient	ρ	-0.4 to 0.4

Table 2: Parameter Values for MWD.

The shortage costs come from a study commissioned by the San Diego County Water Authority, an MWD member, and conducted by the economic consulting firm CIC Research, Inc. The study examines the consequences of incremental water supply cuts of 20%, 60%, and 80% over a two or six-month period, as might occur during a supply interruption caused by an earthquake, or, possibly, a severe drought.[10] The study assessed the costs associated with a two-month, or six-month supply disruption, at varying levels of cutbacks in deliveries (20%, 60%, and 80%). The estimated shortage cost, per af, is $5,885, more than ten times

[10]San Diego County Water Authority (Dec. 28, 1993, Revised Dec. 28, 1994) "The Economic Impact on San Diego County of Three Levels of Water Delivery: 80, 60, or 40 Percent Occurring for Two Months or Six Months," By CIC Research, Inc.

the estimated unit cost of water. As MWD does not currently use an estimate of shortage costs to inform their rate setting, we varied shortage costs in our simulation, using values of \$0, \$5,885, and \$2,500. The total capacity is as reported by MWD in their 2005 Regional Urban Water Management Plan. Table 2 above reports these parameter values.

As discussed earlier, the time intervals, t, in the model should be selected in accord with the actual schedule for decisions by the intermediary, whether it be daily, weekly, quarterly, or annually. MWD sets rates annually, and, hence, we have chosen an annual time step: $t = 1$ corresponds to year one. [11] The question of what time horizon to use is a somewhat trickier issue. We adopt a 10-year horizon, finding, as have Federgruen and Heching (1999), that the year one policy, with this horizon, is stable.

One of the questions of interest is what policy changes are dictated by changes in the relative holding, or storage, and shortage costs. Were storage costs to decrease, or increase, how would the optimal order-up-to-level and price respond? Likewise, if shortage costs nearly double, we are interested as to how the intermediary should respond.

(a) (b)

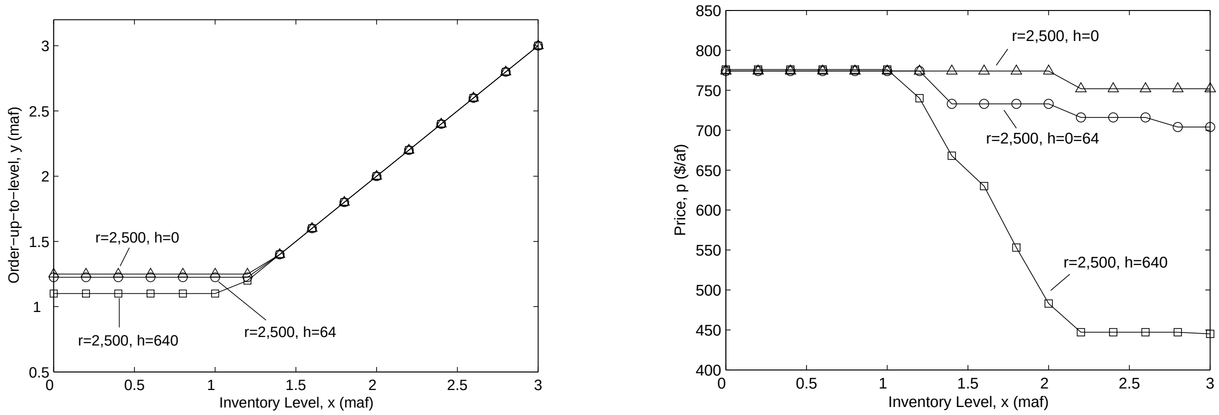

Optimal Order-up-to-level Policy, y^* Optimal Price Policy, p^*

Figure 4: The optimal order-up-to-level, y^* decreases with increasing holding costs, as do the price discounts offered when $x > y^*$.

Holding shortage costs constant at their mid-value, \$2,500/af, we see in Figure 4 that as holding costs increase, the optimal order-up-to-level, y^*, decreases. The optimal list price p^* remains constant, even for a 10 fold increase in the holding cost; however, we see that, for $x > y^*$, the price discounts are greater in magnitude and increase more rapidly the

[11] Decisions about contracts, or order levels, are made on an ongoing basis at MWD, so the reduction of the decision to a single annual review represents a simplification in this respect.

higher the holding cost. We see similar behavior, in Figure 5, where $r = \$0$; the difference in this case is that with no shortage costs, the optimal order-up-to-levels are lower and the price discounting is even more dramatic.

(a) (b)

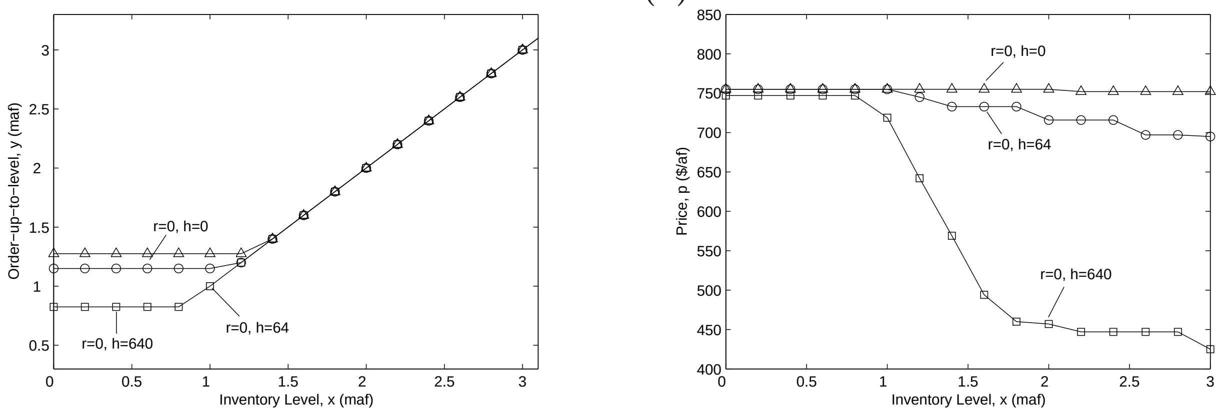

Optimal Order-up-to-level Policy, y^* Optimal Price Policy, p^*

Figure 5: Again, we see that y^* is increasing in h, as are the discounts for $x > y^*$.

For varying shortage costs, with holding cost, h, equal to \$0, we observe that the optimal order-up-to-levels remain almost the same; however, prices increase with shortage costs – up to a threshold at least. When the shortage cost jumps from \$2,500/af to \$5,885/af, the price does not increase any further. Overall, we observe that the policy adjustments associated with increases in holding vs. shortage costs manifest themselves quite differently: changes in holding costs dictate changes in y^*, the optimal order-up-to-level, whereas changes in the shortage costs dictate changes in the optimal p^*. (The holding cost changes do impact the price as well, although only in cases where $x > y^*$ and price discounts are offered below the list price.)

We next consider the impact of shifts in uncertainty. Figures 7 to 9 illustrate the shifts in the optimal policies accompanying these shifts in uncertainty.

In Figure 7, we see that in the case of pure demand uncertainty, the optimal y^* and p^* are increasing in the degree of uncertainty, as captured by the variance of the demand uncertainty, σ_ϵ. Similarly, with pure supply uncertainty, y^* and p^* increase with σ_ξ. However, in Figure 9, we see that as the joint supply and demand uncertainty become extremely high – here, a 15-fold increase in the variances of each – the optimal order level y^* drops to zero. The price continues to increase.

The degree of demand and supply uncertainty also impact the payoff level achieved under the

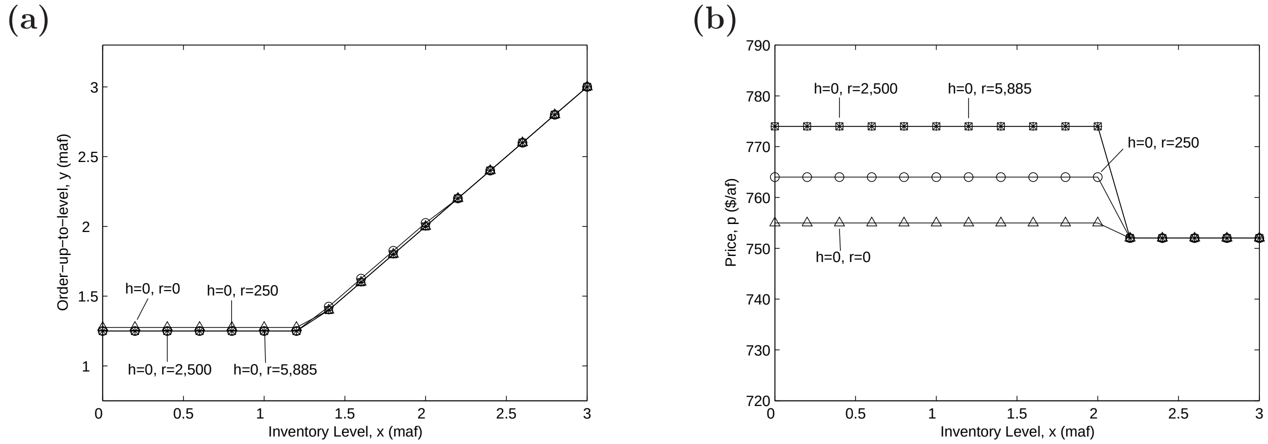

Optimal Order-up-to-level Policy, y^* | Optimal Price Policy, p^*

Figure 6: The optimal y^* remains constant, while the optimal list price, p^* is seen to be increasing with the shortage cost, r.

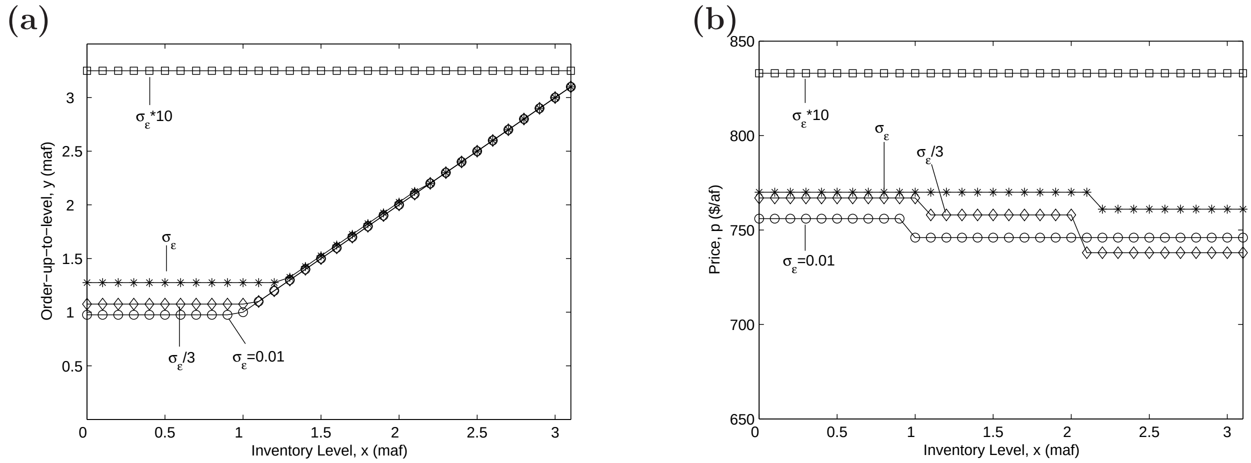

Optimal Order-up-to-level, y^* | Optimal Price, p^*

Figure 7: The optimal order-up-to-level, y^*, and the optimal list price, p^* are increasing in the demand uncertainty, σ_ε. (The supply uncertainty is negligible here, with $\sigma_\xi = 0.01$.) Note that $\rho = 0.4$, $h = \$0$/af, and $r = \$2,500$/af.

optimal policy in period t, Π_t^*. Figures 10, 11, and 12 indicate that for given parameters h and r, some level of uncertainty can actually increase the expected payoff levels. For instance, in the case of pure demand uncertainty, we see the expected payoff level in the first period increasing with σ_ϵ, for $h = \$0$. However, when the storage, or holding, cost increases to $h = \$640$, the payoff levels are seen to be decreasing in the demand uncertainty. In the

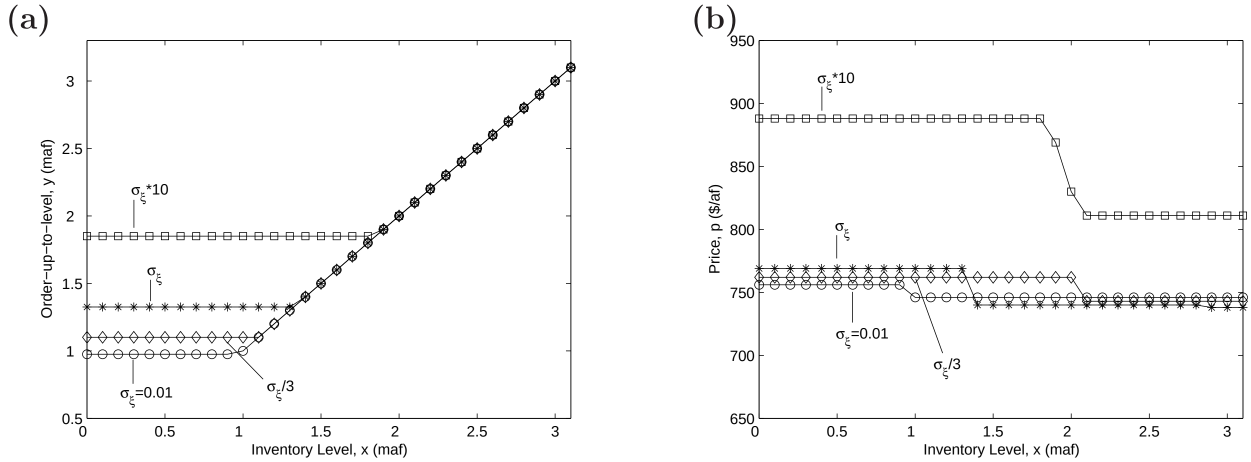

Optimal Order-up-to-level, y^* Optimal Price, p^*

Figure 8: In the case of pure supply uncertainty ($\sigma_\epsilon = 0.01$), both y^* and p^* are seen to be increasing with the degree of uncertainty, σ_ξ. Note that $\rho = 0.4$, $h = \$0$/af, and $r = \$2,500$/af.

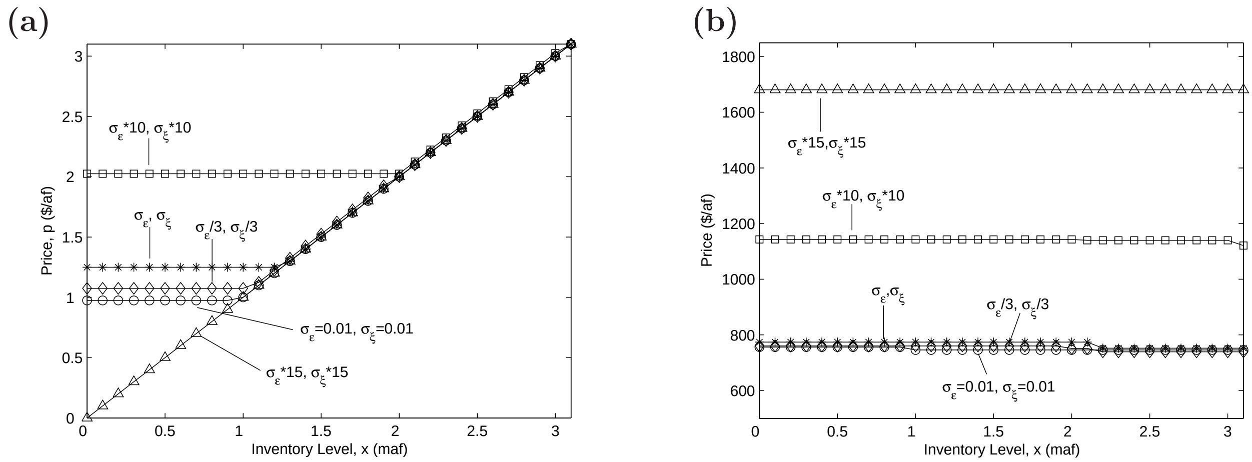

Optimal Order-up-to-level, y^* Optimal Price, p^*

Figure 9: In the case of extremely high demand and supply uncertainty, the optimal order-up-to-level, y^* falls to zero; optimal price, p^* is seen to be increasing in the joint uncertainty. Note that $\rho = 0.4$, $h = \$0$/af, and $r = \$2,500$/af.

case of negligible storage costs, the intermediary can order high levels and set high prices, confining demand to zero in most cases and making large profits (by charging a price almost double that charged for the case of the estimated actual demand uncertainty) in the event

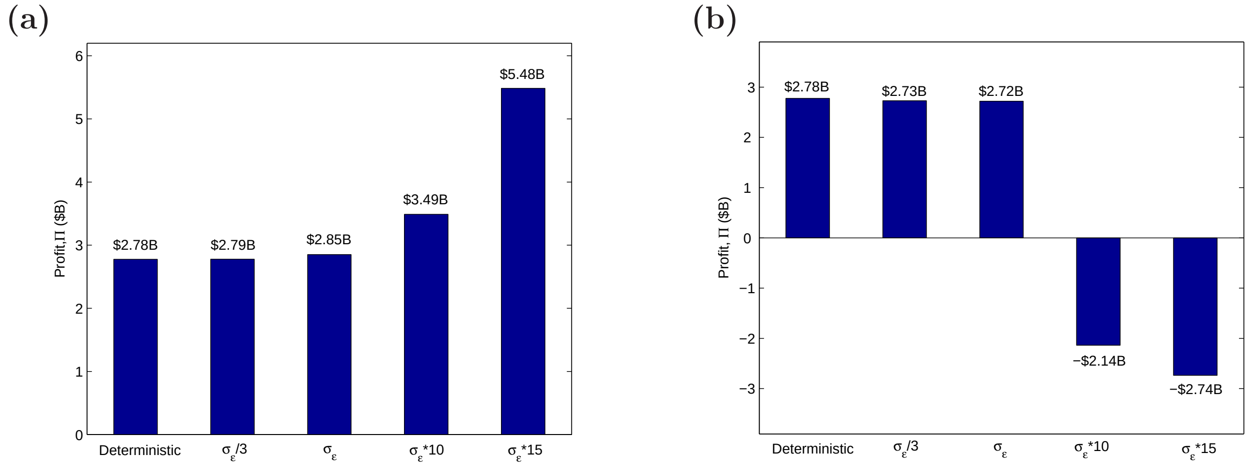

Profit Levels in Period 1, Π_1 Profit Levels in Period 1, Π_1

Figure 10: We observe, in (a), payoff levels increasing in the demand uncertainty, with $h = 0$; however, in case (b), with $h = 640$, payoff levels are *decreasing* in the demand uncertainty.

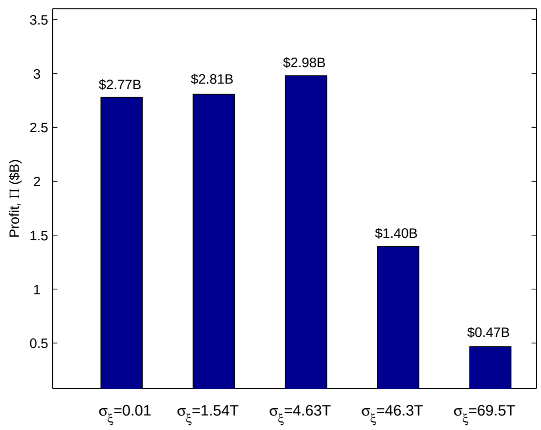

Profit Levels in Period 1, Π_1

Figure 11: In the case of pure supply uncertainty, payoffs initially increase with increasing supply uncertainty, but then decline quite rapidly. Here $\rho = 0.4$, $h = \$0/\text{af}$, and $r = \$2,500/\text{af}$.

of high demand shocks. In the case of high storage costs, it is not possible to cost-effectively store inventory for the occasional chance to meet high demand at a high price; hence we see, in Figure 10 (b) that the expected payoff level is decreasing in the demand uncertainty.

In the case of pure supply uncertainty, with $h = \$0$/af and $r = \$2{,}500$/af the payoff levels increase with the introduction of some supply uncertainty, but then decrease rapidly as the supply uncertainty continues to increase. Presumably this can be attributed to the fact that under such scenarios the intermediary is facing both decreased expected revenues, since expected demand served must be decreasing, and *increasing* expected shortage costs.

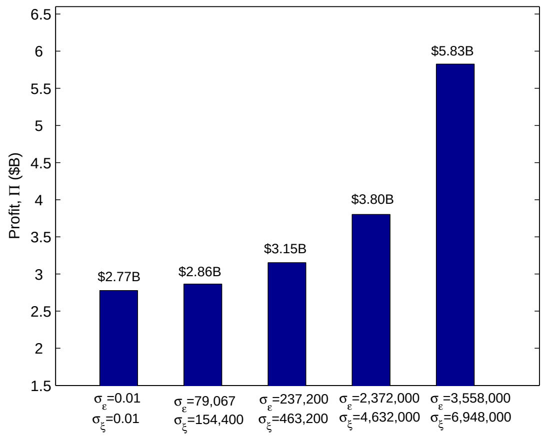

Profit Levels in Period 1, Π_1 -

Figure 12: Profits increase with in supply and demand uncertainty ($\rho = 0.4$).

The third parameter characterizing the uncertainty, apart from the variances of the supply and demand uncertainties, is the correlation coefficient. As described, MWD's current supply, from SWP, is positively correlated with demand. However, the correlation with the true supply, or the source, is negative. In Figure 13, we see that a shift in the correlation from positive to negative dictates a shift in the optimal policy. Namely, the optimal order level, y^* increases. Although the optimal list price, p^*, remains unchanged, we see that the discounting policy in place for the cases where $x > y^*$ is impacted: for higher correlation between supply and demand, price discounts are offered earlier, i.e., at a lower inventory level, x. Note that we see similar behavior in Figure 15, for welfare vs. profit maximization.

Welfare maximization results in a higher order-up-to-level, y^*, and a lower price, p^*. The price shift is significant: under welfare maximization, the price drops by \$300/af, or nearly half, to be closer to \$400/af.

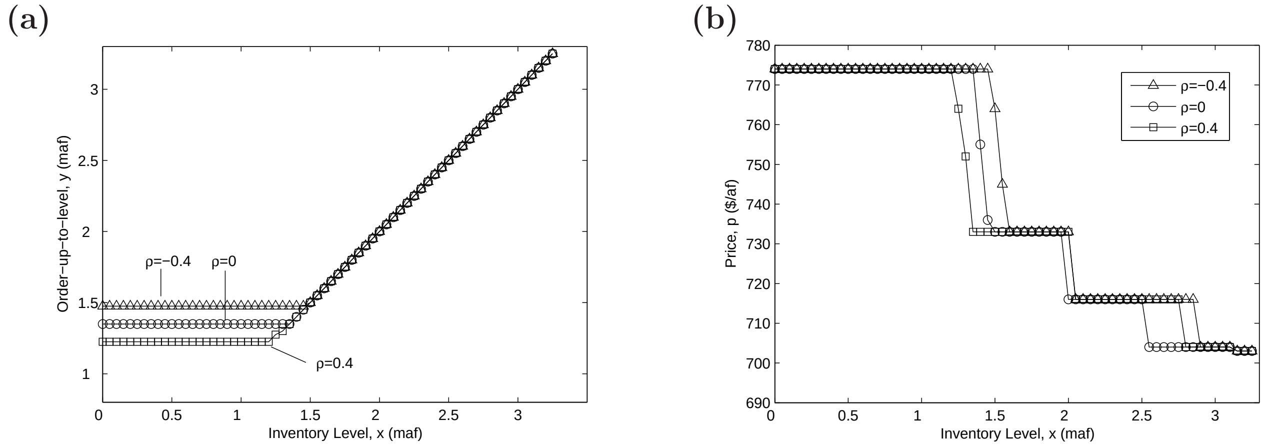

Optimal Order-up-to-level, y^*

Optimal price, p^*

Figure 13: The optimal order-up-to-level, y^*, is decreasing in ρ, the correlation coefficient; the optimal list price, p^* does not shift significantly with ρ. However, note that the price discount in the case $y > x^*$ is offered earlier, i.e., at a lower x, the higher ρ.

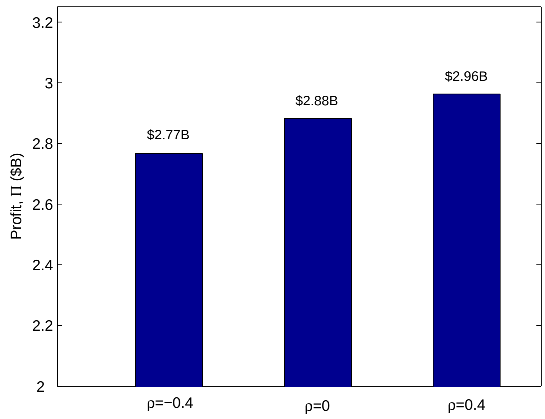

Profit Levels in Period 1, Π_1

Figure 14: Profit levels increase in the correlation coefficient, ρ.

5 Conclusion

The extension of the joint inventory and pricing control model to a setting with two-sided uncertainty, no backlogging, and limits on both the supplier's capacity and the intermediary's maximum order quantity, provides insight into the optimal dynamic pricing and ordering policy of an intermediary attempting to coordinate uncertain supply and uncertain demand,

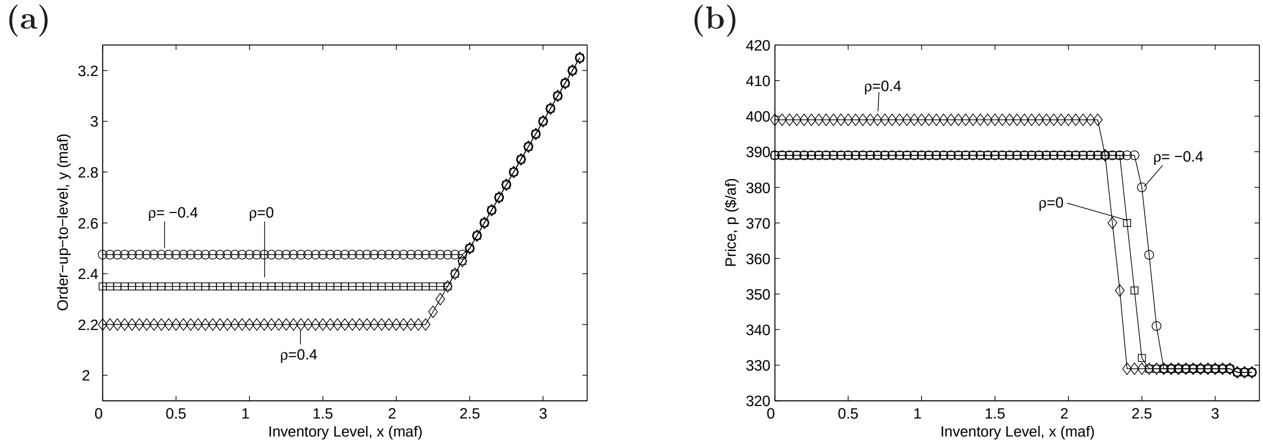

Optimal Order-up-to-level, y^* Optimal price, p^*

Figure 15: The shifts are in accord with those seen under profit maximization; y^* is decreasing in ρ, p^* is unchanged, and the discounting policy shifts to provide earlier discounts (at a lower inventory level, x, the higher ρ.

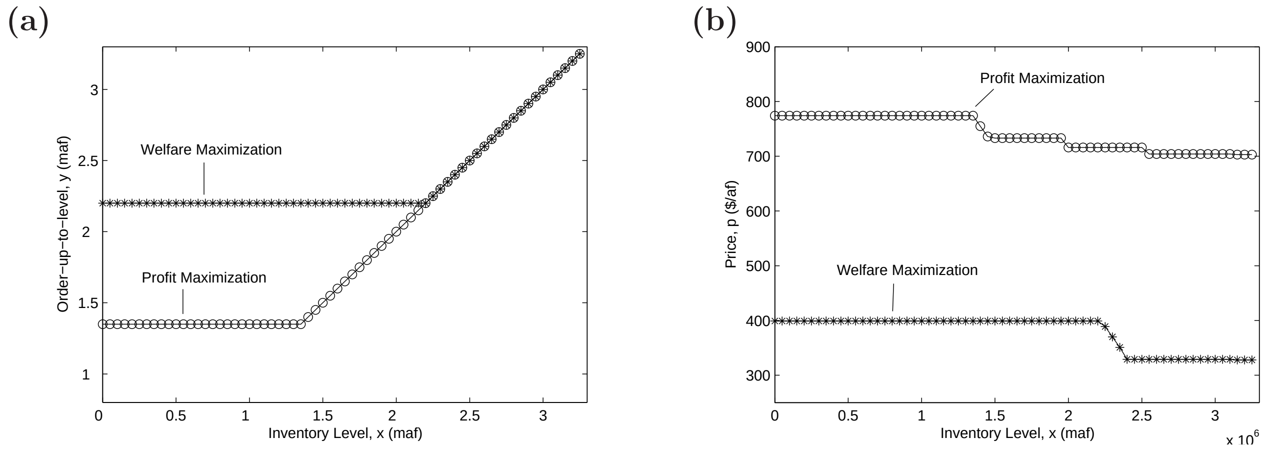

Optimal Order-up-to-level, y^* Optimal price, p^*

Figure 16: As we would expect, the optimal policy under welfare maximization, vs. profit maximization, introduces a higher optimal order-up-to-level, y^*, and a lower price, p^*.

for profit or welfare maximizing purposes. In keeping with earlier results, we find that the optimal policy to be adopted by an intermediary in such cases is a base-stock-list-price policy, or (s, S, P)-policy. This policy entails a fixed list price in the case where the constraint that the order-up-to-level not be less than the inventory state x, i.e., $y^* \geq x$, is non-binding. Federgruen and Heching (1999) find that in the case where that constraint is in fact binding, so that the intermediary does not order anything, it never optimal to increase price, i.e., $p^*(x) \leq P$. We show that this conclusion does not extend to our setting; in cases where demand is highly inelastic and the intermediary is profit maximizing, it is best to not decrease price with the available inventory to take advantage of monopoly rationing.

We expand the traditional application of inventory and pricing control models to consider allocative efficiency (in terms of a welfare-maximizing policy), in addition to profit maximization, finding that, somewhat unsurprisingly, the optimal ordering level is higher under welfare maximization, while the price is generally lower. The application of inventory and pricing control models to public pricing problems has not, to date, received significant attention. The advantages of such applications over existing public pricing models include the ability to capture storage dynamics and to model ordering decisions with fixed short-run capacity. As noted in the introduction, the extant public pricing literature provides models that treat joint pricing and *capacity* decisions, not pricing and *ordering* decisions, thereby ignoring storage. Our public pricing model for water may also be applicable to other natural resources, such as oil and natural gas, and even to electricity, where the storability of electricity – via dams, fly wheels, and fuel cells – call for a framework other than the traditional peak-load pricing models.

Finally, our case study of the Metropolitan Water District of Southern California (MWD) illustrates the potential of such models to inform current policy with respect to the following five aspects.

1. *Holding costs*: as holding costs increase, the intermediary should increase its order-up-to-level, y^*. The impact on price less pronounced: when the constraint $y^* \geq x$ is binding, we see that price discounting is greater, and more rapid, in cases where h is higher. The optimal list price does not shift under increasing holding costs.

2. *Shortage costs*: In the case of increasing shortage costs, the intermediary should rely on price adjustments, versus order quantity adjustments, to maintain optimality. This is in contrast to the case of increasing holding costs and indicates the value of the *dual* policy levers price, and order-up-to-level.

3. *Demand and supply uncertainty*: in the cases of pure demand uncertainty, pure supply uncertainty, and two-sided uncertainty, we see that increases in uncertainty imply increases in y^* and p^*, up to a point. In the case of extremely high (15-fold above the estimated actual variance) supply uncertainty, we see that the optimal order-up-to-level is driven to zero, while the price continues to increase (insuring that in all cases

except those with the highest demand shocks, demand will be zero). In the case of pure demand uncertainty, the impact of additional demand uncertainty on the payoff function is tied to the storage, or holding, cost, h. We provide an example where, for negligible holding costs, the intermediaries payoff function is increasing in the demand uncertainty, σ_ϵ. However, the result is reversed for a higher storage cost ($h = \$640$), in which case we see that the payoff function is actually decreasing in σ_ϵ. In the case of pure supply uncertainty, the payoff function increase as with σ_ξ initially, but it is clear that as σ_ξ continues to increase and the intermediary faces increasing expected shortage costs and decreasing expected revenues, the payoff function is decreasing. These results imply the benefit of cheaper storage in the face of demand uncertainty, and the overall desirability of some level of supply uncertainty in the contractual agreements entered into by the intermediary.

4. *Correlated demand and supply uncertainty*: the correlation of supply and demand impacts the optimal order-up-to-level, y^*, under both profit and welfare maximization; y^* is seen to be decreasing in ρ. However, a shift in p^* is only visible in the welfare maximization case, in which case p^* is decreasing in ρ.

5. *Pricing under supply windfalls*: in contrast to current practices by MWD, wherein a single annual rate is set regardless of supply levels, the inventory and pricing control model dictates a price adjustment in cases where the available inventory exceeds the optimal order-up-to-level. In most cases the price will be adjusted downwards, with the exception of profit maximization in the face of inelastic demand, in which case price is adjusted upwards.

Appendix A: Proofs

Proof of Proposition 1. By definition, the expected demand served is

$$Q(p,y) = \int_{\mathbb{R}} \left(\int_{-\infty}^{\min\{\lambda y+\xi,\bar{x}\}-(\alpha-\beta p)} (\alpha - \beta p + \varepsilon)\varphi(\varepsilon,\xi)d\varepsilon + \int_{\min\{\lambda y+\xi,\bar{x}\}-(\alpha-\beta p)}^{\infty} \min\{\lambda y + \xi, \bar{x}\}\varphi(\varepsilon,\xi)d\varepsilon \right) d\xi.$$

Hence, by differentiation (using Leibniz rule) we obtain that Q is increasing in price p,

$$Q_p(p,y) = -\beta P(\tilde{d}_t(p) \leq \tilde{s}_t(y)) < 0,$$

and that Q is increasing in the order-up-to level y,

$$Q_y(p,y) = \lambda P(\tilde{s}_t(y) < \min\{\tilde{d}_t(p), \bar{x}\}) > 0,$$

so that the cross-derivative of Q with respect to (p,y) becomes

$$Q_{py}(p,y) = -\beta\lambda \int_{-\infty}^{\bar{x}-\lambda y} \varphi(\lambda y + \xi - (\alpha - \beta p), \xi)d\xi < 0.$$

The last inequality implies the submodularity of Q with respect to (p,y). We now consider the consumer surplus $\mathrm{CS}(p,y)$, which can be written in the form

$$\mathrm{CS}(p,y) = E\left[\left(\frac{\alpha}{\beta} - p \right) \tilde{q}(p,y) + \frac{\tilde{\varepsilon}\tilde{q}(p,y)}{\beta} - \frac{\tilde{q}^2(p,y)}{2\beta} \right].$$

Hence, the derivative of expected consumer surplus with respect to p becomes

$$\mathrm{CS}_p(p,y) = \frac{\partial}{\partial p} E\left[\left(\frac{\alpha}{\beta} - p \right) \tilde{q}(p,y) \right] + \frac{\partial}{\partial p} E\left[\frac{\tilde{\varepsilon}\tilde{q}(p,y)}{\beta} \right] - \frac{\partial}{\partial p} E\left[\frac{\tilde{q}^2(p,y)}{2\beta} \right].$$

The second term can be computed as follows:

$$\begin{aligned} \frac{\partial}{\partial p} E\left[\frac{\tilde{\varepsilon}\tilde{q}(p,y)}{\beta} \right] &= \frac{1}{\beta}\frac{\partial}{\partial p} \left(\int_{\mathbb{R}} \left[\int_{-\infty}^{\min\{\lambda y+\xi,\bar{x}\}-(\alpha-\beta p)} \varepsilon(\alpha - \beta p + \varepsilon)\varphi(\varepsilon,\xi)d\varepsilon \right.\right. \\ &\qquad \left.\left. + \int_{\min\{\lambda y+\xi,\bar{x}\}-(\alpha-\beta p)}^{\infty} \varepsilon \min\{\lambda y + \xi, \bar{x}\}\varphi(\varepsilon,\xi)d\varepsilon \right] d\xi \right) \\ &= -\int_{\mathbb{R}} \int_{-\infty}^{\min\{\lambda y+\xi,\bar{x}\}-(\alpha-\beta p)} \varepsilon\varphi(\varepsilon,\xi)d\varepsilon d\xi. \end{aligned}$$

Similarly, the third term simplifies to

$$\begin{aligned} -\frac{\partial}{\partial p} E\left[\frac{\tilde{q}^2(p,y)}{2\beta} \right] &= -\frac{1}{2\beta}\frac{\partial}{\partial p} \left(\int_{\mathbb{R}} \left[\int_{-\infty}^{\min\{\lambda y+\xi,\bar{x}\}-(\alpha-\beta p)} (\alpha - \beta p + \varepsilon)^2\varphi(\varepsilon,\xi)d\varepsilon \right.\right. \\ &\qquad \left.\left. + \int_{\min\{\lambda y+\xi,\bar{x}\}-(\alpha-\beta p)}^{\infty} (\min\{\lambda y + \xi, \bar{x}\})^2\varphi(\varepsilon,\xi \right] d\varepsilon \right) d\xi \\ &= \int_{\mathbb{R}} \int_{-\infty}^{\min\{\lambda y+\xi,\bar{x}\}-(\alpha-\beta p)} (\alpha - \beta p + \varepsilon)\varphi(\varepsilon,\xi)d\varepsilon d\xi, \end{aligned}$$

so that

$$\begin{aligned}\mathrm{CS}_p(p,y) &= \left(\frac{\alpha}{\beta}-p\right)Q_p - Q - \int_{\mathbb{R}}\int_{-\infty}^{\min\{\lambda y+\xi,\bar{x}\}-(\alpha-\beta p)} \varepsilon\varphi(\varepsilon,\xi)d\varepsilon d\xi \\ &\quad + \int_{\mathbb{R}}\int_{-\infty}^{\min\{\lambda y+\xi,\bar{x}\}-(\alpha-\beta p)} (\alpha-\beta p+\varepsilon)\varphi(\varepsilon,\xi)d\varepsilon d\xi.\end{aligned}$$

Using the expression for Q_p obtained earlier, we notice that the first term on the right-hand side cancels the last two terms, whence

$$\mathrm{CS}_p(p,y) = -Q(p,y) < 0.$$

Similarly, we obtain for the derivative of the expected consumer surplus with respect to the order-up-to level y that

$$\begin{aligned}\mathrm{CS}_y(p,y) &= \left(\frac{\alpha}{\beta}-p\right)Q_y + \frac{\lambda}{\beta}\int_{-\infty}^{\bar{x}-\lambda y}\int_{-\infty}^{\min\{\lambda y+\xi,\bar{x}\}-(\alpha-\beta p)} \varepsilon\varphi(\varepsilon,\xi)d\varepsilon d\xi \\ &\quad -\frac{\lambda}{\beta}\int_{-\infty}^{\bar{x}-\lambda y}\int_{\min\{\lambda y+\xi,\bar{x}\}-(\alpha-\beta p)}^{\infty} \min\{\lambda y+\xi,\bar{x}\}\varphi(\varepsilon,\xi)d\varepsilon d\xi,\end{aligned}$$

or, substituting our earlier expression for Q_y, equivalently,

$$\mathrm{CS}_y(p,y) = \frac{\lambda}{\beta}E\left[\tilde{d}-\tilde{s}\,\middle|\,\tilde{s}<\min\{\tilde{d},\bar{x}\}\right]P(\tilde{s}<\min\{\tilde{d},\bar{x}\}) > 0.$$

By differentiating $\mathrm{CS}_p(p,y)$ with respect to the order-up-to level y we obtain that

$$\mathrm{CS}_{py}(p,y) = -Q_y < 0,$$

since $Q_y > 0$. Hence, the expected consumer surplus is submodular with respect to (p,y), which completes our proof. ■

Proof of Proposition 2. Let $(p,y),(\hat{p},\hat{y}) \in \mathbb{R}\times X$ with $(p,y) \leq (\hat{p},\hat{y})$. For $x=\bar{x}$ it is $F(x|p,y) = F(x|\hat{p},\hat{y}) = 1$. For any $x \in [0,\bar{x})$ it is

$$F(x|p,y) - F(x|\hat{p},\hat{y}) = \int_{\mathbb{R}}\left(\int_{\min\{\lambda y+\xi,\bar{x}\}-(\alpha-\beta p)-x}^{\min\{\lambda\hat{y}+\xi,\bar{x}\}-(\alpha-\beta\hat{p})-x} \varphi(\varepsilon,\xi)d\varepsilon\right)d\xi \geq 0,$$

so that indeed $F(\cdot|p,y) \preceq_{\mathrm{FOSD}} F(\cdot|\hat{p},\hat{y})$. ■

Proof of Proposition 3. We note first that for any $t\in\{0,\ldots,T-1\}$ the state-transition probability $F(x_{t+1}|p_t,y_t)$ is independent of the current state x_t. Hence, the expected payoff in the next period, $E\left[\Pi_{t+1}(\tilde{x}_{t+1})|\,p_t,y_t\right]$, is independent of x_t. The function

$$\Pi_t(x) - cx = \max_{(p,y)\in\mathbb{R}\times[x,\bar{x}]}\left\{g(p,y) + \delta E\left[\Pi_{t+1}(\tilde{x}_{t+1})|\,p,y\right]\right\}$$

must be nonincreasing in x, for an increase in x only further constrains choice set without influencing the maximand. The preceding argument remains unchanged for the last period, when $t = T$, by setting $\Pi_{T+1}(x) \equiv 0$. ■

Proof of Proposition 4. Let $(p, y), (\hat{p}, \hat{y}) \in \mathbb{R} \times X$ with $(p, y) \leq (\hat{p}, \hat{y})$. By Proposition 3, the function $\Pi_t(x) + cx$ is decreasing on X. Hence, using the fact that by Proposition 2 it is $F(\cdot|p, y) \preceq_{\text{FOSD}} F(\cdot|\hat{p}, \hat{y})$ we obtain that

$$E\left[\Pi_t(\tilde{x}_{t+1}) - c\tilde{x}_{t+1} | \, \hat{p}, \hat{y}\right] \leq E\left[\Pi_t(\tilde{x}_{t+1}) - c\tilde{x}_{t+1} | \, p, y\right],$$

which, by the linearity of the expectation operator, implies (1). ■

Proof of Proposition 6. (i) Consider the case of welfare maximization, where $\omega = 1$. Then

$$g_{py} = Q_y + pQ_y - L_{py} + \text{CS}_{py},$$

which, using the expression for $\text{CS}_{py} = -Q_y$, derived in the proof of Proposition 1, implies that

$$g_{yp} = pQ_y - L_{py} < 0,$$

i.e., the intermediary's per-period payoff is submodular in (p, y). (ii) Consider now the case of profit maximization, where $\omega = 0$. Then the nonnegative term $Q_y = \lambda P(\tilde{s}_t(y) < \min\{\tilde{d}_t(p), \bar{x}\})$ is not necessarily cancelled out by $-L_{py} = -\beta\lambda(r + h) \int_{-\infty}^{\bar{x}-\lambda y} \varphi(\lambda y + \xi - (\alpha - \beta p), \xi) d\xi < 0$, especially if β is small. This completes our proof. ■

Proof of Proposition 5. Let $x \in X$. The proof proceeds in two steps.

Step 1. We show that for any $y \in [x, \bar{x}]$ the per-period payoff $g(p, y)$ is coercive in p. For this, consider

$$g_p(p, y) = Q(p, y) + pQ_p(p, y) - L_p(p, y) + \omega \, \text{CS}_p(p, y) =$$

$$Q(p, y)(1 - \omega) - \beta \left(p(1 - \omega) + \omega \frac{\alpha}{\beta} + (r + h) \right) P(\tilde{d}_t(p) \leq \tilde{s}_t(y)) + \beta r -$$

$$\omega \left(\int_{\mathbb{R}} \int_{-\infty}^{\min\{\lambda y + \xi, \bar{x}\} - (\alpha - \beta p)} \varepsilon \varphi(\varepsilon, \xi) d\varepsilon d\xi - \int_{\mathbb{R}} \int_{-\infty}^{\min\{\lambda y + \xi, \bar{x}\} - (\alpha - \beta p)} (\alpha - \beta p + \varepsilon) \varphi(\varepsilon, \xi) d\varepsilon d\xi \right)$$

$$= Q(p, y)(1 - \omega) - \beta \left(p(1 - \omega) + \omega \frac{\alpha}{\beta} + (r + h) \right) P(\tilde{d}_t(p) \leq \tilde{s}_t(y)) + \beta r - \omega \zeta$$

where $\zeta = \left(E[\tilde{\varepsilon}|\tilde{d}_t(p) < \tilde{s}_t(y)] - E[\tilde{q}|\tilde{d}_t(p) < \tilde{s}_t(y)] \right) P(\tilde{d}_t(p) < \tilde{s}_t(y))$. Since

$$\lim_{p\to\infty} P(\tilde{d}_t(p) \le \tilde{s}_t(y)) = \lim_{p\to\infty} P(\tilde{d}_t(p) \le \tilde{s}_t(\bar{x})) = 1,$$

for any $\eta \in (0,1)$ there is a $\hat{p}_\eta$, independent of y, such that $p \ge \hat{p}_\eta$ implies that $P(\tilde{d}_t(p) \le \tilde{s}_t(y)) > 1 - \eta$. Hence, by setting $\eta = 1/2$ for any $\rho > 0$, we have the following:

$$g_p(p,y) \le \bar{x} - \beta \left(p(1-\omega) + \omega \frac{\alpha}{\beta} + (r+h) \right) (1-\eta) + \beta r - \omega\zeta < -\rho\beta/2 < 0 \qquad (2)$$

for all $p \ge \max\{\hat{p}(x), r + h - \omega \left(\frac{\alpha}{\beta} - \frac{2\zeta}{\beta}\right) + \rho + \frac{2\bar{x}}{\beta}\}$, where $\hat{p}(x)$ is such that $P(\tilde{d}_t(\hat{p}(x)) \le \tilde{s}_t(x)) \ge 1/2$. The upper bound $\bar{p}(x)$ obtains by taking the limit for $\rho \to 0^+$. The upper bound $\bar{p}(x)$ is nonincreasing in the state x, since $P(\tilde{d}_t(p) \le \tilde{s}_t(x))$ is nondecreasing in x for any given p. We therefore conclude that increasing p beyond $\bar{p}(x)$ can only decrease expected payoffs. Let us now consider $p = 0$. For any $y \in [x, \bar{x}]$ it is

$$g_p(0,y) = Q(0,y)(1-\omega) - \beta \left(\omega \frac{\alpha + \tilde{\varepsilon}}{\beta} + (r+h) \right) P(\tilde{d}_t(p) \le \tilde{s}_t(y)) + \beta r + \omega\zeta > 0,$$

so that a price of zero cannot be optimal: increasing p yields a strict improvement in expected payoffs. We have thus shown that at time $t = T$, any price p_T that maximizes per-period rewards must lie in $[0, \bar{p}(x)]$.

Step 2. We now show by induction that for any $t \in \{1, \dots, T-1\}$ the expected value at time $t+1$ is strictly decreasing in p if only p is large enough. Assume that this is true for $\bar{\Pi}_{t+1}(p,y)$. Then

$$\bar{\Pi}_t(p,y) = g(p,y) + cE\left[\tilde{x}_{t+1}|\, p, y\right] + \delta\bar{\Pi}_{t+1}(p,y)$$

is decreasing in p for large enough p if $g(p,y) + cE\left[\tilde{x}_{t+1}|\, p, y\right]$ is decreasing in p for large enough p. We have that

$$\frac{\partial}{\partial p} E\left[\tilde{x}_{t+1}|\, p, y\right] = \beta P(\tilde{d}_t(p) \le \tilde{s}_t(y)) \le \beta,$$

so that we can conclude from (2) with $\rho = (2 + \hat{\rho})c > 0$ for any $\hat{\rho} > 0$ that

$$g_p(p,y) + \frac{\partial}{\partial p} E\left[\tilde{x}_{t+1}|\, p, y\right] \le c\beta - c(2+\hat{\rho})\beta/2 = -\hat{\rho}\beta/2 < 0$$

for all $p \ge \max\{\hat{p}(x), r + h - \omega \left(\frac{\alpha}{\beta} + \frac{2\zeta}{\beta}\right) + \rho + \frac{2\bar{x}}{\beta}\}$, where $\hat{p}(x)$ is such that $P(\tilde{d}_t(\hat{p}(x)) \le \tilde{s}_t(x)) \ge 1/2$. As in Step 1, the upper bound $\bar{p}(x)$ obtains by taking the limit for $\hat{\rho} \to 0^+$, which concludes our proof. ■

Proof of Proposition 7. The optimal price depends on x if and only if the optimal order-up-to level is equal to x. (i) Under welfare maximization, when $\omega = 1$, the proof follows immediately from the submodularity of g of Proposition 6 by backward induction, starting for $t = T$. (ii) Under profit maximization, when $\omega = 0$, part (ii) of Proposition 6 states that per-period $g(p, y)$ can be supermodular, so that for small enough discount factors δ the intermediary's total discounted payoff is then also supermodular in (p, y), which together with Milgrom and Shannon's (1994) monotonicity theorem concludes our proof. ■

Proof of Proposition 8. We have that $y \in \{x, \bar{x}\}$. Let y^* denote the optimal order-up-to-level, y. Now suppose that $y^* < x$, in which case y^* is infeasible. Let $y^{*'}$ denote the optimal order-up-to-level when the constraint that $y^* \geq x$ is binding. Suppose that $y^{*'} = x' > x$. By Proposition 4, $E[\Pi_t(\tilde{x}_{t+1})|p, \hat{x'}] \leq E[\Pi_t(\tilde{x}_{t+1})|p, \hat{x}]$, which contradicts $y^{*'} = x' > x$. Hence, we conclude, $y_t^*(x) \in \{x, \bar{y}_t^*\}$.

By Proposition 5 and the coerciveness of $\Pi_t(p, y)$, we have established the existence of an interior maximizer, p^*: it is neither optimal to set $p = 0$, nor is it optimal to charge ∞. It is left to establish that p^* is independent of x whenever the constraint it non-binding, the definition of a list-price-policy. This independence is evident through examination of the derivatives of $g(p, y)$ above. ■

Appendix B: Supplementary Material

Table 5 below details MWD's rationing policy under the Incremental Interruption and Conservation Program (IICP), which has since been replaced by the Water Surplus and Drought Management Plan (WSDM), adopted in 1999. Stage I of the IICP was adopted in December of 1990. Then, in January of 1991, a Stage III policy was put in place. In February 1991, MWD declared a Stage IV, which remained in place until March 1992, when it downgraded to a Stage III, finally returning to Stage I in April of 1992.

In 2002 MWD moved to an "unbundled" rate system, whereby rates are presented as the sum of the costs of providing different services. (Previous to 2004 MWD had a lumped rate, setting a price for, say, "full service (non-interruptible) treated" water or "untreated interim agricultural" water, without delineating the different cost components of providing this water.) Tier 1 Supply Rate represents the current cost of acquiring one unit (one af) of

Stage	Reduction in Non-Firm Deliveries	Conservation Goal for Firm Deliveries
Stage I	Voluntary	Goal of 10%
Stage II	20%	5%
Stage III	30%	10%
Stage IV	40%	15%
Stage V	50%	20%

Table 3: MWD's Incremental Interruption and Conservation Plan, Stages I through IV.

water; the Tier 2 Supply Rate represents the cost of augmenting the existing water supply, by an af. The System Access Rate is the cost of infrastructure for water deliveries. The Water Stewardship Rate goes towards environmental protection. The Treatment Surcharge is the cost of treating an af of water. Finally, the last two charges are fixed fees, where the Readiness-to-Serve Charge (RTS) is designed to cover the cost of providing emergency storage services, in case of supply disruption in the event of an earthquake or of a severe drought. The $80M fee is divided amongst the 26 member agencies based on the rolling 10-year average of deliveries to each agency; that is, if an agency's 10-year average of deliveries represents 10% of the total average 10-year deliveries, then that agency will cover 10% of the RTS. Finally, the Capacity Charge for each agency is levied on the maximum daily flow to the agency within the last three years, between May 1 and September 20. This charge is designed such that it "recovers the cost of providing peak capacity within the distribution system." In our model, we consider only the price full service treated water (Tier 1 and Tier 2 water with all of the additional charges, i.e., with the system access rate, water stewardship rate, system power rate, and treatment surcharge added to the Tier 1 and Tier 2 base fee.)

Table 5 reports the rates MWD has charged over the last twenty one years - through the end of 2006 - for full service treated water, which is what we model. As noted in Table 5 above, MWD offers different types of water, at different rates, including untreated full service water, treated and untreated "interim agricultural" water, and treated and untreated"seasonal storage" water. The full service water represents the largest portion of MWD deliveries and is considered "non-interruptible," whereas the interim agricultural and seasonal storage supplies are deemed "interruptible." In 2003, MWD introduced a block pricing system, with "Tier 1" and "Tier 2" charges for water consumed above a certain level. That is, all member agencies were provided the opportunity to enter into long-term (10-year) contracts with MWD for delivery of up to 90% of their historical water use at the Tier 1 rate. Usage above this amount would be charged the Tier 2 rate. The Tier 2 rate is denoted in parentheses,

Description of Charges	Rate for FY 2006
Tier 1 Supply Rate ($/af)	$73
Tier 2 Supply Rate ($/af)	$169
System Access Rate ($/af)	$152
Water Stewardship Rate ($/af)	$25
System Power Rate ($/af)	$81
Treatment Surcharge ($/af)	$122
Readiness-to-Serve Charge ($M)	$80M
Capacity Charge ($/cfs)	$6,800

Table 4: Unbundled Water Rates for Financial Year 2006.

after the Tier 1 rate, for years 2003–2006 above.

Year(s)	Full Service Treated Rate
1986–1990	$230/af
1991	$261/af
1992	$322/af
1993	$385/af
1994	$412/af
1995-1996	$426/af
1997-2002	$431/af
2003	$408/af ($489/af)
2004	$418/af ($499/af)
2005	$443/af ($524/af)
2006	$453/af ($549/af)

Table 5: Water Rates Over the Past 21 Years.

Table 5 documents the estimated shortage costs by sector (per af of undelivered water), from the CIC Research, Inc., study, “The Economic Impact on San Diego County of Three Levels of Water Delivery: 80, 60, or 40 Percent Occurring for Two Months or Six Months.” The total costs per undelivered af of water sum to $5,885.

Sector	Shortage Cost
Agriculture	\$569/af
O.R.P.	\$123/af
Construction	\$449/af
Trans. & P.U.	\$255/af
Manufacturing	\$763/af
Wholesale Trade	\$149/af
F.I.R.E.	\$1,082/af
Retail Trade	\$575/af
Services	\$1,858/af
Government	\$61.3/af
Total	≈ \$5,885/af

Table 6: Estimated shortage costs in 2006 for undelivered water, by sector. One af cut in supplies is assumed to be uniformly applied across all sectors. ("F.I.R.E" stands for "Financial Services and Real Estate," "Trans. and P.U." is Transportation and Public Utilities, and "O.R.P." is Other Resource Production.)

References

[1] Arrow, K., Harris, T., Marschak, J. (1951) "Optimal Inventory Policy," *Econometrica*, Vol. 19, No. 3, pp. 250–272.

[2] Bös, D. (1985) "Public Sector Pricing," in: Auerbach, A.J., Feldstein, M. (Eds.), *Handbook of Public Economics, Vol. I,* Elsevier, Amsterdam, NL, pp. 128–211.

[3] Chen, X., Simchi-Levi, D. (2004) "Coordinating Inventory Control and Pricing Strategies with Random Demand and Fixed Ordering Cost: The Finite Horizon Case," *Operations Research*, Vol. 52, No. 6, pp. 887–896.

[4] Chen, F.Y., Ray, S., Song, Y. (2006) "Optimal Pricing and Inventory Control Policy in Periodic-Review Systems with Fixed Ordering Cost and Lost Sales," *Naval Research Logistics*, Vol. 53, No. 2, pp. 117–136.

[5] Clark R.M., Stevie, G. (1981) "A Water Supply Cost Model Incorporating Spatial Variables," *Land Economics*, Vol. 57, No. 1, pp. 18–32.

[6] Crew, M.A., Fernando, C., Kleindorfer, P. (1995) "The Theory of Peak Load Pricing: A Survey," *Journal of Regulatory Economics*, Vol. 8, No. 3, pp. 215–248.

[7] Draper, A.J., Jenkins, M.W., Kirby, K.W., Lund, J.R., Howitt, R.E. (2003) "Economic-Engineering Optimization for California Water Management," *Journal of Water Resources Planning and Management*, Vol. 129, No. 3, pp. 155–164.

[8] Draper, A.J., Lund, J.R. (2004) "Optimal Hedging and Carryover Storage Value," *Journal of Water Resources Planning and Management*, Vol. 130, No. 1, pp. 83–87.

[9] Elmaghraby, W., Keskinocak, P. (2003) "Dynamic Pricing in the Presence of Inventory Considerations: Research Overview, Current Practices, and Future Directions," *Management Science*, Vol. 49, No. 10, pp. 1287–1309.

[10] Federgruen, A., Heching A. (1999) "Combined Pricing and Inventory Control Under Uncertainty," *Operations Research*, Vol. 47, No. 3, pp. 454–475.

[11] Gallego G., Van Ryzin, G. (1994) "Optimal Dynamic Pricing of Inventories with Stochastic Demand over Finite Horizons," *Management Science*, Vol. 40, No. 8, pp. 999–1020.

[12] Howitt, R.E., Vaux, H.J. (1984) "Managing Water Scarcity: An Evaluation of Interregional Transfers," *Water Resources Research*, Vol. 20, No. 7, pp. 785–792.

[13] Kim, D., Lee, W.J. (1998) "Optimal Joint Pricing and Lot Sizing with Fixed and Variable Capacity," *European Journal of Operational Research*, Vol. 109, No. 1, pp. 212–227.

[14] Lee, W.J. (1994) "Optimal Order Quantities and Prices with Storage Space and Inventory Investment Limitations," *Computers and Industrial Engineering*, Vol. 26, No. 3, pp. 481–488.

[15] Mills, E.S. (1959) "Uncertainty and Price Theory," *Quarterly Journal of Economics*, Vol. 73, No. 1, pp. 116–130.

[16] Petruzzi, N.C., Dada, M. (1999) "Pricing and the Newsvendor Problem: A Review with Extensions," *Operations Research*, Vol. 47, No. 2, pp. 183–194.

[17] Polatoglu, L.H. (1991) "Optimal Order Quantity and Pricing Decisions in Single-Period Inventory Systems," *International Journal of Production Economics*, Vol. 23, No. 1–3, pp. 175–185.

[18] Rajan, A., Steinberg, R., Steinberg, R. (1952) "Dynamic Pricing and Ordering Decisions by a Monopolist," *Management Science*, Vol. 38, No. 2, pp. 240–262.

[19] Thowsen, G.T. (1975) "A Dynamic Non-Stationary Inventory Problem for a Price/Quantity Setting Firm," *Naval Research Logistics Quarterly*, Vol. 22, pp. 461–476.

[20] Whitin T.M. (1953) *The Theory of Inventory Management*, Princeton University Press, Princeton, NJ.

[21] Whitin T.M. (1955) "Inventory Control and Price Theory," *Management Science*, Vol. 2, No. 1, pp. 61–68.

[22] Zabel, E. (1972) "Multiperiod Monopoly Under Uncertainty," *Journal of Economic Theory*, Vol. 5, No. 3, pp. 524–536.