Designing New Water Rates for Los Angeles

W. Michael Hanemann
University of California, Berkeley

Introduction

By April 1991, the customary end of the rainy season in California, it became evident that the fourth year of a drought would bring severe disruption to the California water system and that, for the first time in fourteen years, urban users in southern California would face cutbacks in their water supply. In Los Angeles, the Department of Water and Power (LADWP) announced that it was asking for a 15% reduction in water use systemwide. In the next couple of months, citizens responded eagerly and water use fell. So did water revenues. The reduction in revenues provoked a financial crisis. LADWP had to return to its customers and announce a rate increase. This created a political furor, to which Mayor Bradley responded by appointing a Blue Ribbon Committee on Water Rates with 24 members, including citizens from all geographic areas and major constituencies in the city. One of the committee members was an economist, Professor Darwin Hall. The Committee and its staff set about various investigations, including a detailed study of the marginal costs of water supply. In March 1992, as that study was nearing completion, the Committee asked me and a Berkeley colleague, Professor Shmuel Oren, to serve as Technical Advisers. Working closely with Committee members and staff, I devised a new rate structure that the Committee adopted in its Report of June, 1992. With modifications described below, the rate structure was approved by the City Council in January, 1993 and went into effect the next month.

In this article, I want to describe some of the issues that arose, and the lessons that I learned, in the course of my work with the Committee.

Contrasts Between Water and Power

In recent years, there has been an a growing tendency to point to the electricity supply industry as a role model for how the urban water industry should evolve with respect to issues such as pricing, integrated resource planning, and demand side management. Moreover, water utilities themselves have begun to make comparisons between the existing degree of reliability set by regulators for electric utilities and the lower levels of reliability currently being attained in the water industry. Clearly, there are important similarities between the two industries. But there are also differences, some of which may turn out to have considerable economic significance. Hence, it is useful to begin by reviewing where the two industries are alike and where different.

From an economic point of view, one key fact stands out about the urban water industry both nationally and in California. This is that — to a much greater degree than electricity and other public utilities — the water industry is publicly owned as opposed to investor-owned. Nationally, approximately 85% of the public water supply systems serving more than 3,300 customers are publicly owned. In California, publicly owned water systems serve over 90% of the state's population. The major consequence of this public ownership is that, from an economic (as opposed to public health) perspective, the water utility industry is essentially an unregulated industry. In this respect, it differs substantially from the electricity industry which, in California, has been closely regulated by both the PUC and the California Energy Commission.

A second fundamental characteristic of the water supply industry is its capital intensity. For the water industry nationally, the asset requirement per dollar of revenue (i.e. the ratio of capital assets to annual revenue) has been estimated at about $10-12; this is 3 to 4 times the capital intensity of the telephone and electric utility industries, and 5 to 6 times that of the railroad industry. Because of its extraordinary capital intensity, investment and pricing decisions are of crucial importance to the industry's efficiency. Given the combination of public ownership and absence of economic regulatory oversight, it is possible that there may have been some shortcomings in this regard.

Perhaps the most important differences from an economic point of view have to do with storage and transportation. Compared to electricity, water is relatively easier to store but harder to transport (wheel). This has several implications. With water it is possible — and indeed common - to have terminal storage capable of holding an amount equal to several months' usage, even one or two years' usage. Therefore, there is little chance of short-run outages caused by sudden spikes in demand. On the other hand, water is expensive to transport and cannot be moved at short notice — it is not like electricity which can be purchased on a spot market from some location as much as 1,000 miles away. When the reservoir runs dry, it will stay dry until the rains come, or a new source of supply is negotiated. These differences mean that the costs of attaining any given degree of reliability are different for water than for electricity, so that what constitutes the optimal degree of reliability for one may not be optimal for the other.

There are also implications for marginal cost pricing. The existence of spot markets for electricity together with the presence of techniques for generating electricity such as gas turbine or hydropower than can essentially be turned on and off at will imply that the marginal cost curve for electricity is fundamentally elastic, as shown in Figure 1A: although it will be more expensive than infra-marginal sources, you can always be sure of obtaining an incremental supply of electricity from somewhere. By contrast, I would argue that the marginal cost curve of water for an urban water utility is more apt to be inelastic, as shown in Figure 1B: you can obtain water from particular sources that have been arranged ahead of time, but once you exhaust these sources, you are out of luck — you would have to turn to some form of rationing or voluntary demand reduction.

The difference between the shapes of the marginal cost curves in Figures 1A and 1B matters because it implies a different approach to pricing. In the case of the elastic marginal cost curve in Figure 1A, the challenge in ratesetting is essentially to determine which are the marginal sources of supply and what are their costs. Once these are known, the rate structure follows fairly simply — you want to ensure that users at the margin pay whatever is the cost of the marginal source. By contrast, in the case of the inelastic marginal cost curve of Figure 1B, the task may be different — the focus may be chiefly on determining what price it will take to ration demand to keep it within the bounds of available supply. In these circumstances, while it is still useful to know about marginal cost, it is even more important to know about the demand function. In short, you only need an engineer in order to cope with the situation in Figure 1A, while in Figure 1B you need a social scientist.

Key Features of M&I Demand

Implicitly, at least, much of the existing economic theorizing about water pricing tends to assume a single, homogeneous, temporally invariant demand for water. This overlooks two features of M&I demand that are of considerable importance — climate-induced temporal variability, and heterogeneity in the population of water users.

Figure 2, reproduced from the LADWP publication "Statistical Report for the Fiscal years 1981 -1990," shows clearly the cyclical nature of M&I demand. In any year, demand in the peak month is 50-70% larger than the minimum month's demand. This cyclical pattern of demand is a classic example of the type of situation that calls for peak load pricing. To be sure, this does not necessarily apply to the costs of water supply and major storage: precisely because there is storage, the marginal cost of supply is the same year-round. But, it certainly does apply to transmission, treatment, and distribution costs: one would think that the sizing of these types of capacity is driven by peak demand and therefore, according to the conventional economic formula, these capacity costs should be borne by peak users: they should be recovered entirely through charges on peak use, not off-peak use.

Beside the variation in demand from month to month, and year to year, due to the changes in climate, there is also significant variation in use among different customers within the same general class. For example, Table 1 shows data on the size distribution of average monthly water use across single-family residential accounts in the LADWP service area in 1988 (i.e., before the drought). Obviously, one factor that influences the usage per account is the size of the household — but I don't believe that this is the only factor at work here. In an average month in 1988, usage per residential account varied from as little as 25 gallons per day to as high as 22,400 gallons per day (there were 5 accounts using this amount). According to the literature from AWWA and similar sources, a typical range for residential indoor use in a single-family unit would be 65-80 gallons per capita per day. The literature suggests that outdoor use — including lawn watering, swimming pools, and cooling — is much more variable (and much more price elastic), and could range from, say, as low as 25 gallons per capita per day to 100 gallons per capita per day, or more. I would assume that most households have 3 to 5 members (the average for the MWD service area is under 3.5 persons per household). I would suspect that much of the variation in Table 1 reflects not household size but differences in patterns of use — primarily differences in outdoor use but also, to a lesser extent, differences in indoor use (for example, whether or not people bother to fix leaky faucets, and whether or not they have remodelled and installed gigantic new bathtubs). In short, these are differences in demand behavior that I believe would not be picked up by most demand functions currently used in the literature. The variation in use does not reflect difference in price, since these customers were all paying the same, flat-rate price. While it surely is influenced by lot-size, a variable commonly included in demand functions, I think it reflects something more than that, namely differences in preferences.

This is supported by the data shown in Figure 3, taken from Kiefer and Dziegielewski (1991), which derive from a survey of 830 single-family residences in 1990 located throughout the MWD service area. The survey attempted to measure both how much water was used for outdoor irrigation and how much should have been used, given the size, location and other characteristics of the lot. In effect, this standardizes for differences in the variables that would shift the demand curve for outdoor water use; the ratio of actual outdoor use to "agronomically correct" use should be unity in every case. In fact, of course, it isn't — it varies from around 10% (i.e., severe underwatering) to about 625% (overwatering). About 55% of the households underwater their yards, while 45% overwater. I would guess that this variation reflects primarily differences in ability, knowledge, carefulness (in the sense of Leibenstein's (1976) X-inefficiency), or preferences — factors which are conventionally omitted from consideration in the economic literature on the demand for water.

Implications for Water Rates

In my view, the presence of this type of heterogeneity has important implications for the design of a water rate schedule because it suggests an alternative to what I would consider the conventional strategy. The conventional strategy focuses on the average water use per customer, rather than the entire distribution, and hopes to change this average via prices or other instruments which, in effect, shift the entire distribution to the left. The alternative strategy considers the entire distribution and attempts to change the average by changing the shape of the distribution — i.e., by shifting some parts of the distribution such as the right tail much more than other parts.

How, one might ask, is this consistent with conventional economic theory ? The answer is that, while it is unusual, it can indeed be explained in terms of an economic model, albeit a model that is somewhat broader than the conventional one. Specifically, one can tell at least two stories which support creating a price differential rather than simply raising a flat-rate price. The first has to do with information and psychology. Creating a price differential may convey information to consumers by suggesting that there is some norm attached to the usage level at which the switch in rates occurs — the implied message from the authorities is "we think you should be able to get by using no more than this amount, although we are leaving you free to consume more if you wish, albeit at a higher price." For this to work, the switchpoint has to be located at a level which is consistent with this normative overtone. In the LADWP case, we considered switchpoints that corresponded to 175% or 200% of the median household use.

The second story has to do with variation in the demand elasticity at different levels of use. One could distinguish, for example, between "core" uses and "discretionary" uses, with the former being considerably less elastic than the latter. Indoor use for toilets, bathing, cooking and washing may be something that the residential customer is not readily inclined to change, given the appliances that the household owns. Outdoor use is far more price elastic. Certain types of appliance choice may also be quite price-elastic — since about 1990, for example, Los Angeles and Santa Monica have offered subsidies for retrofitting ultra-low flush toilets, with considerable success [see Chesnutt et al. (1992)]. The idea, therefore, is that one has different prices targeted at different end uses or at different types of demand, with the higher price focused on the more price-responsive elements.

One other consideration should be mentioned. Water is an unusual commodity in that, while its use can in principle be controlled with a high degree of precision — you can turn off the faucet whenever you want — in practice most consumers have no idea, either ex ante or ex post, of exactly how much they are using. They know what type of dishwasher they own, for example, but they don't know how much water it consumes. Therefore, there is likely to be a significant either random or unplanned component in the household's demand for water. The conventional view in economic theory is that one should focus on moving the consumer from one point on the demand curve to another via higher prices, not on shifting the demand curve via the provision of information, exhortation, etc. For water, it is less obvious that these other strategies should be shunned. If this is accepted, it reinforces several of the points made above — for example, that the informational content of a price system is important, and that rate setting is about influencing behavior. It also seriously undermines the conventional welfare-theoretic argument for marginal cost pricing, since there are now two different sets of preferences, one before the information campaign and the other after. Whether a pricing strategy that deviates from strict marginal cost pricing principles is efficient or not depends on whether one uses ex ante or ex post preferences to assess welfare — if ex post preferences are used, the deviation from marginal cost pricing may still be efficient.

Application to LADWP Rates

I come, at last, to the ratesetting exercise at LADWP. We followed four principles. First, we wanted users at the margin to face the marginal cost of water supply, which we calculated taking reclaimed wastewater as the marginal source. Second, I wanted there to be a seasonal price differential reflecting what we believed to be the seasonal differential in marginal cost, following the argument made above about the peak users bearing the capacity costs of treatment, transmission and distribution. In 1985, LADWP had introduced seasonally differentiated flat-rates; the summer rate was initially set at about 15% higher than the winter rate and by 1992 the differential had risen to 25%. I felt, however, that this differential was too small to attract much notice from consumers, as well as too small to reflect the real differences in correctly calculated peak and off-peak marginal costs.

Thirdly, I wanted to incorporate a two-tier price structure aimed at shrinking the right-hand tail of the distribution of demand; this meant having a substantial price differential between the two blocks and locating the switchpoint at a level of usage where demand might be reasonably responsive to the price incentive. Fourthly, I believed that, for the price incentive to work, it was not necessary to make everybody actually pay the higher price on some units of their consumption. I felt that the incentive would still be effective for consumers who were below the switchpoint, as long as it was sufficiently close that the higher price loomed in their consciousness and could influence their purchase of water-using appliances.

Based on my analysis of the limited data available, I felt that the switchpoint for a single-family residential account should probably be located somewhere in the range of 400-600 gallons/account/day. This assumes 4 or 5 persons in the household, an indoor usage of 70-80 gallons/capita/day, and some basic outdoor usage (recognizing that the outdoor usage, and possibly the indoor usage, does not increase linearly with the number of persons in the household). I am assuming, of course, that these are the "core" uses which are substantially less elastic than the other uses — although, in the absence of solid data, this is a hunch rather than a scientific certainty.

This contrasts with most of the increasing block rate schedules currently existing in other California cities and elsewhere. Some of those rate structures have three, four or five blocks, often with quite small differentials between the blocks. This excessively dilutes the incentive effects. I prefer the simpler but more powerful signal associated with two blocks and sharper price differentials. Moreover, most of the existing rate structures locate their switchpoints at 300 gallons/account/day or lower. In my view, this is too low to elicit a substantial response — I am willing to trade off having a smaller fraction of customers actually in the higher block for the prospect of a greater behavioral response in which a larger proportion of users in that block switch out of it.

The rates that the Blue Ribbon Committee recommended are shown in Table 2. An important feature is that there are two sets of rate schedules — one for normal years and another for drought years. The idea behind the latter is to set down, ahead of time, the principles that will be followed when it comes to adjusting water rates in the course of a drought situation. Focusing first on the normal rates, there are two basic schemes — one for single-family residences and the other for the remaining customer classes. Under the rate structure proposed by the Committee, there is one rate for use up to 525 gallons/account/day and another, higher, rate for consumption in excess of this amount. For other customers — multi-family residential, commercial and industrial — there is a single rate for consumption in winter; in the summer, this same rate applies for consumption up to 125% of winter consumption, while a higher rate applies for consumption beyond this level. The second block rate is the same for all classes of user and reflects an estimate of LADWP's marginal cost of supply. It differs between summer and winter to allow for the seasonal differences in marginal cost. The rate for the first block varies among customer classes and is set so as to meet the revenue targets for that class. These rates should be compared with the rates instituted in 1992, after the drought crisis had passed but before the Committee had developed its own recommendations, which were a flat rate for all customer classes alike amounting to $1.55/ccf in winter and $1.76/ccf in summer.

In drought years, the same type of structure still applies, but it is modified in two ways to adjust to the shortage situation. First, the switchpoints at which the second block commences are reduced, roughly in proportion to the severity of the shortfall. Second, the rate charged in this second block is raised to equal what the Committee estimated to be the rationing price that would equilibrate demand to supply, given the shortfall. Thus, the Committee estimated the equilibrium price needed to accomplish a 15% cutback within the LADWP service area shortage at $4.44/ccf, and similarly with the prices for the other shortage levels of shortage. These equilibrium prices were based on an analysis of the actual experience in the second half of 1991, when tiered prices were introduced as a means of both raising revenue and rationing via prices, following the crisis described at the beginning of this paper. Since they involve an extrapolation from limited data, they should be viewed as approximations.

Six months after the Blue Ribbon Committee presented its report, the Los Angeles City Council adopted a rate ordinance, of which the Normal Year section is shown in Table 3. This follows the Committee's recommendations quite closely. The main change was to raise the switchpoint for single-family residential accounts and differentiate it by season, placing it at 575 gallons/account/day in winter and 725 gallons/account/day in summer. In addition, the high block price was raised slightly, and the low block price was reduced substantially in the light of revised revenue requirement calculations for the LADWP system. The drought year rates were modified correspondingly, preserving the same basic structure as the Committee had recommended.

In the future, I hope to work with LADWP staff to monitor the impact of the new rate structure on water usage in its service area. We hope to track a sample of individual accounts from the different customer classes in order to quantify the effects of the new rate structure and, more generally, to test the assumptions that underlie it. Hopefully, this will improve our understanding of demand behavior and allow us in the future to substitute informed analysis for educated guesses.

References

Chesnutt, Thomas W., Anil Bamezai, and Casey McSpadden. 1992. Mapping the Conserving Effect of Ultra Low Flush Toilets: Implications for Planning. Report prepared for Metropolitan Water District by A&N Technical Services, Inc. Santa Monica.

Kiefer, Jack C., and Benedykt Dziegielewski. 1991. Analysis of Residential Landscape Irrigation in Southern California. Report prepared for Metropolitan Water District by Planning and Management Consultants, Ltd, Carbondale, IL.

Leibenstein, H. 1976. Beyond Economic Management. Cambridge, Mass: Harvard University Press.

City of Los Angeles. 1992. Mayor's Blue Ribbon Committee on Water Rates, A Consensus Approach to Water Rates.

Note

I want to express my deep gratitude to Darwin Hall and to Gerald Gewe, Richard West and their staff at LADWP for their kindness and helpfulness throughout.

W. Michael Hanemann is in the Department of Agricultural and Resource Economics at the University of California, Berkeley.

Tables 1-3 and Figures 1-3 follow.

TABLE 1: AVERAGE WATER USE IN 1988 BY LADWP RESIDENTIAL CUSTOMERS

ACCOUNT USAGE (gal/day)	# OF CUSTOMERS	CUMULATIVE PERCENT OF CUSTOMERS	CUMULATIVE PERCENT OF CONSUMPTION
0 - 75	13,381	3.3	0.4
75 - 150	31,041	11.0	2.6
150 - 200	32,740	19.1	6.0
200 - 250	39,100	28.8	11.1
250 - 300	41,172	39.0	17.7
300 - 350	39,832	48.9	25.1
350 - 400	35,474	57.7	32.7
400 - 450	30,251	65.2	40.0
450 - 500	25,102	71.4	46.8
500 - 550	20,422	76.5	52.9
550 - 600	16,406	80.5	58.2
600 - 650	13,239	83.8	62.8
650 - 700	10,803	86.5	66.9
700 - 800	15,475	90.3	73.4
800 - 900	10,308	92.9	78.3
900 - 1000	7,341	94.7	82.2
1000 - 1100	5,178	96.0	85.3
1100 - 1200	3,661	96.9	87.6
1200 - 1300	2,708	97.5	89.5
1300 - 1400	2,098	98.1	91.1
1400 - 1500	1,614	98.5	92.4
1500 - 1750	2,477	99.1	94.6
1750 - 2000	1,319	99.4	96.0
2000 - 2500	1,296	99.7	97.6
> 2500	1,086	100.0	100.0
TOTAL	403,524		

TABLE 2: WATER RATES PROPOSED BY LADWP BLUE RIBBON COMMITTEE

NORMAL YEAR RATES

	PRICE IN LOW BLOCK ($/CCF)	SWITCH POINT	PRICE IN HIGH BLOCK ($/CCF) WINTER	PRICE IN HIGH BLOCK ($/CCF) SUMMER
RESIDENTIAL				
Single-Family	$1.71	525 gallons/day	2.27	$2.92
Multi-Family	$1.71	125% of winter use	NA	$2.92
NON-RESIDENTIAL	$1.78	125% of winter use	NA	$2.92

DROUGHT YEAR RATES

	PRICE IN LOW BLOCK ($/CCF)	SWITCH POINT	PRICE IN HIGH BLOCK ($/CCF)
RESIDENTIAL			
Single-Family			
10% Shortage	$1.71	475 gallons/day	$3.70
15% Shortage	$1.71	450 gallons/day	$4.44
20% Shortage	$1.71	425 gallons/day	$5.18
25% Shortage	$1.71	400 gallons/day	$6.05
Multi-Family			
10% Shortage	$1.71	115% of adjusted winter use	$3.70
15% Shortage	$1.71	115% of adjusted winter use	$4.44
20% Shortage	$1.71	110% of adjusted winter use	$5.18
25% Shortage	$1.71	110% of adjusted winter use	$6.05
NON-RESIDENTIAL			
10% Shortage	$1.78	115% of adjusted winter use	$3.70
15% Shortage	$1.78	115% of adjusted winter use	$4.44
20% Shortage	$1.78	110% of adjusted winter use	$5.18
25% Shortage	$1.78	110% of adjusted winter use	$6.05

TABLE 3: NORMAL YEAR WATER RATES ADOPTED BY LA CITY COUNCIL

	PRICE IN LOW BLOCK ($/CCF)	SWITCH POINT	PRICE IN HIGH BLOCK ($/CCF) WINTER	SUMMER
RESIDENTIAL				
Single-Family	$1.14	WINTER: 575 gallons/day	$2.33	
		SUMMER: 725 gallons/day		$2.98
Multi-Family	$1.14	125% of winter use	NA	$2.98
NON-RESIDENTIAL	$1.21	125% of winter use	NA	$2.98

FIGURE 1:
THE SHAPE OF THE MARGINAL COST CURVE FOR URBAN WATER

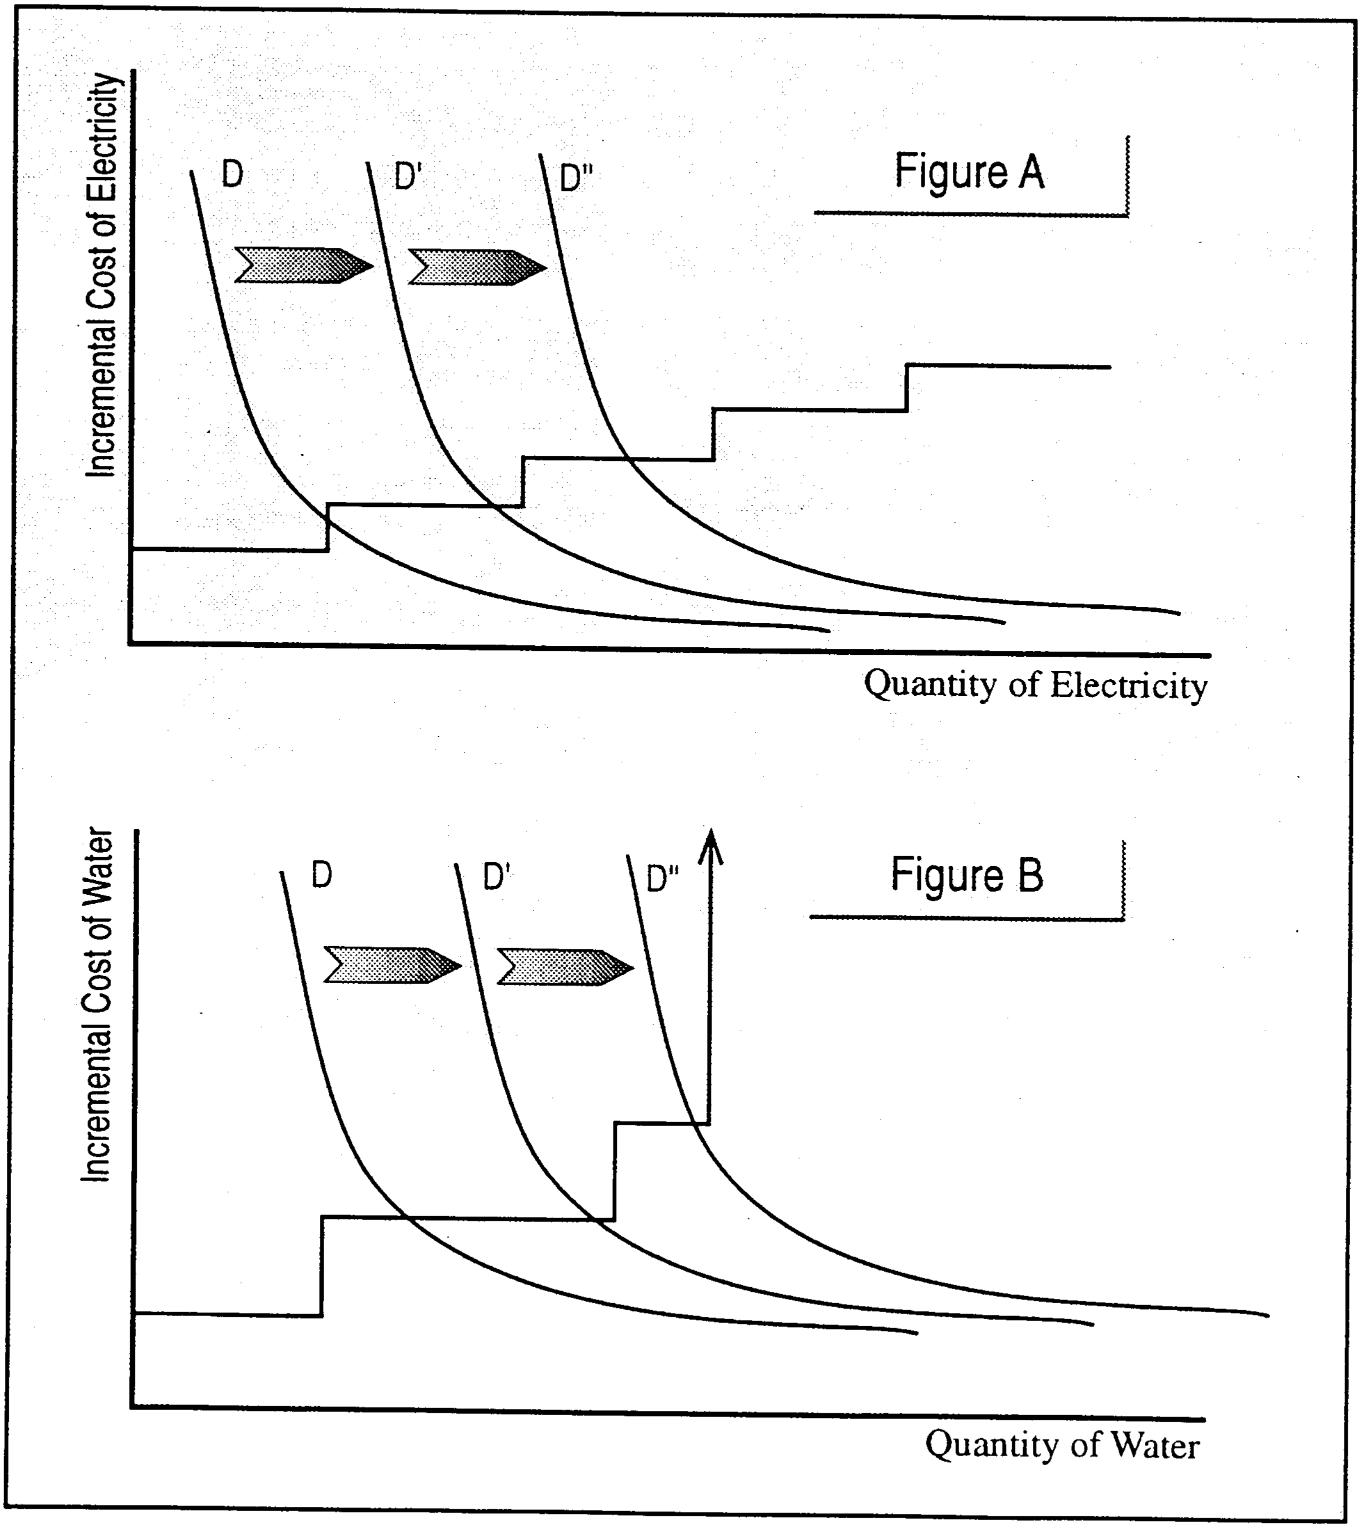

FIGURE 2:

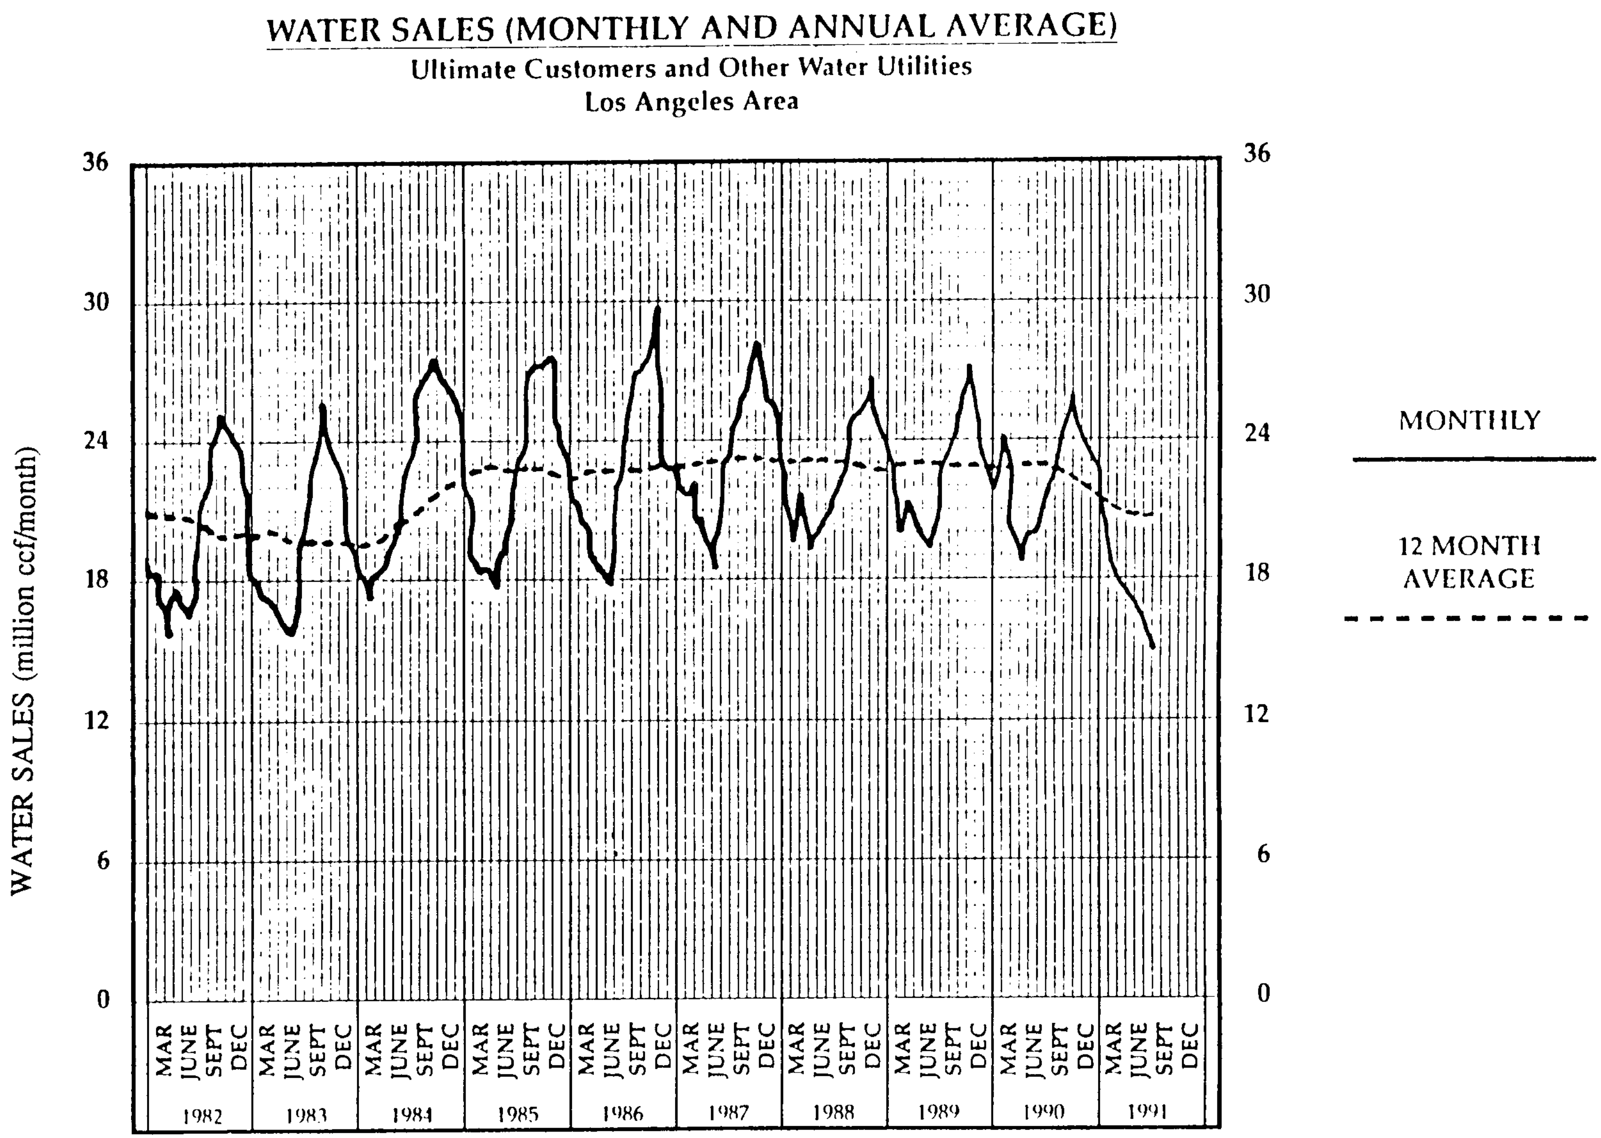

WATER SALES (MONTHLY AND ANNUAL AVERAGE)
Ultimate Customers and Other Water Utilities
Los Angeles Area
WATER SALES (million ccf/month)
36
30
24
18
12
6
0
MAR
JUNE
SEPT
DEC
1982
1983
1984
1985
1986
1987
1988
1989
1990
1991
MONTHLY
12 MONTH AVERAGE

FIGURE 3:
DEFICIT/SURPLUS IRRIGATION BY HOUSEHOLDS IN MWD SERVICE AREA (Nov 1989 - Oct 1990)