The Value of Scarce Water: Measuring the Inefficiency of Municipal Regulations*

April 24, 2009

Erin T. Mansur
Yale School of Management
135 Prospect Street
PO Box 208200
New Haven, CT 06520-8200

Sheila M. Olmstead
Yale School of Forestry
and Environmental Studies
195 Prospect Street
New Haven, CT 06511
Tel. (203) 432-6247
Fax (203) 436-9150
sheila.olmstead@yale.edu
Corresponding author

* Mansur is an Associate Professor, School of Management, Yale University. erin.mansur@yale.edu. Olmstead is an Associate Professor, School of Forestry and Environmental Studies, Yale University. sheila.olmstead@yale.edu. We are grateful for input from Patrick Bayer, Nathaniel Keohane, Sharon Oster, Robert Stavins, Christopher Timmins, and seminar participants at Yale, Harvard, Dartmouth, Columbia, Camp Resources, and the California Occasional Workshop on Environmental Economics. We thank two anonymous referees and our associate editor for comments that improved the quality of the manuscript. The authors, alone, are responsible for any errors.

The Value of Scarce Water:
Measuring the Inefficiency of Municipal Regulations

ABSTRACT

Rather than allowing water prices to reflect scarcity rents during periods of drought-induced excess demand, policy makers have mandated command-and-control approaches, like the curtailment of certain uses, primarily outdoor watering. Using unique panel data on residential end-uses of water, we examine the welfare implications of typical drought policies. Using price variation across and within markets, we identify end-use specific price elasticities. Our results suggest that current policies target water uses that households, themselves, are most willing to forgo. Nevertheless, we find that use restrictions have costly welfare implications, primarily due to household heterogeneity in willingness-to-pay for scarce water.

***Journal of Economic Literature* Classifications:** L510, Q580, Q250, Q280, L950

Keywords: Resource Allocation, Market-based Regulation, Residential Water Demand, Drought Policy

1. Introduction

Where markets are not employed to allocate scarce resources, the potential welfare gains from a market-based approach can be estimated. We assess the potential welfare gains and possible distributional outcomes from using prices to reduce urban water consumption, rather than water rationing, during periods of drought. Between January 1997 and June 2007, moderate to extreme drought conditions affected, on average, 20 percent of the contiguous United States (National Climatic Data Center 2007).[1] During droughts, municipal water restrictions focus almost exclusively on the residential sector (which comprises one-half to two-thirds of urban consumption) rather than on commercial and industrial water users. Rather than allowing prices to reflect scarcity rents during periods of excess demand, policy makers have mandated the curtailment of certain uses, primarily outdoor watering, requiring the same limitations on consumption of all households. If indoor demand is not perfectly inelastic, or if households are heterogeneous in willingness-to-pay for water under conditions of scarcity, a price-based approach to drought policy has a theoretical welfare advantage over water rationing.

Using unique panel data on residential end-uses of water for 1,082 households in 11 urban areas in the United States and Canada, we examine the implications of the current approach to urban drought. Using price variation across and within markets, we identify price elasticities specific to indoor and outdoor water demand. We find that outdoor watering restrictions do mimic household reactions to price increases, on average, as outdoor demand is much more price elastic than indoor use. However, the real advantage of market-based approaches lies in their accommodation of heterogeneous marginal benefits. During periods of scarcity, in which regulators impose identical frequency of outdoor watering, shadow prices for

[1] This 10-year average masks much variation; in June 2007, 41 percent of the contiguous United States was experiencing moderate to extreme drought (National Climatic Data Center 2007).
[3] We assume that supply, in the short-run situation of drought, is perfectly inelastic.

the marginal unit of water may vary greatly among households. We estimate separate end-use elasticities for four heterogeneous household groups, based on income and lot size, to assess the degree of household heterogeneity in these markets. We estimate shadow prices for each consumer and utility-level market-clearing prices under four drought policy scenarios of increasing stringency. We then simulate the effects of moving to a market based approach, in comparison to a two-day-per-week outdoor watering restriction.

Our results have important implications for municipal policies. The welfare gains from allowing substitution across uses within the average household amount to about \$56 during an average summer drought. Allowing substitution across households, as well as within, welfare gains from a price-based approach are approximately \$92 per household, about 28 percent of average annual household expenditures on water in our sample. These direct welfare gains would also come with potential savings in enforcement and monitoring costs – volumetric metering and billing systems are already in place for water consumption in most North American cities, while command-and-control approaches currently require direct observation of individual households' outdoor consumption.

Switching to a price-based policy would have allocative consequences. Drought prices would enable those customers who are the least price sensitive, the wealthy consumers with large lots, to reduce consumption much less than low-income households with small lots, who may stop using outdoor water entirely, even under modest drought conditions. The distributional consequences of these changes depend on the assignment of water rights. If utilities retain these rights, their profits would rise by an estimated \$136 per customer. Households would be worse off by \$44, on average. Rebates from water suppliers to consumers could make everyone better off.

This paper provides the first estimates of end-use water demand elasticities in the peer-reviewed literature, and the most comprehensive estimates to date of the welfare loss from rationing water during droughts. Rationing approaches are ubiquitous in the water sector, especially in arid regions. For example, during a 1987-1992 drought in California, 65-80 percent of urban water utilities implemented outdoor watering restrictions (Dixon et al. 1996). In 2008, 75 percent of Australians lived in communities with some form of mandatory water use restrictions (Grafton and Ward 2008). Our results suggest that the economic losses from such approaches may be quite substantial.

Note that the analysis does not consider whether water prices, in general, are set at efficient levels – only whether the policy instruments used to constrain use during a drought minimize the costs of achieving consumption reductions of a given magnitude. The market-based approach we simulate would re-allocate consumption during a drought so as to maximize the net benefits of the drought-constrained water supply. In this context, the questions we address are "second-order" questions, if current prices depart significantly from efficient levels.

The paper proceeds as follows. In Section 2, we discuss the theory and related literature. Section 3 presents the econometric models, and Section 4 describes the data. In Section 5, we present our price elasticity estimates for end uses and consumer groups. Section 6 discusses the economic consequences of current regulatory policies and distributional impacts of switching to a price-based allocation mechanism. In Section 7, we conclude.

2. Theory and Related Literature

2.1. Markets and Social Welfare in General

The questions addressed in this paper arise from the general theoretical literature on the conditions under which gains in social welfare are possible through the introduction of markets for managing scarcity (Weitzman 1977, Suen 1990, Newell and Stavins 2003). There are theoretical and empirical estimates of the gains from increasing the influence of markets on traffic congestion on roadways (Small and Yan 2001, Parry and Bento 2002) and at airports (Daniel 2001, Pels and Verhof 2004); electricity generation (Gilbert *et al.* 1996, De Vany and Walls 1999, Kleit and Terrell 2001); and water supply across sectors (Howe *et al.* 1986, Hearne and Easter 1997). The gains from price-based approaches to allocation in these cases derive largely from heterogeneity in consumers' marginal benefits.

A related, prolific literature has focused on welfare comparisons of market-based and command-and-control approaches to environmental policy (Pigou 1920, Crocker 1966, Dales 1968, Montgomery 1972, Baumol and Oates 1988, Tietenberg 1995). The gains from market-based approaches in the case of pollution control policies derive from cost heterogeneity among regulated firms. The advantage of markets over prescriptive policies for pollution control has been demonstrated empirically for the regulation of lead in gasoline (Kerr and Maré 1997), sulfur dioxide in power plant emissions (Burtraw et al. 1998), and many other applications.

In the literature on the economics of pollution control, analysts have focused on distributional outcomes where pollutants are non-uniformly mixed, creating variation in the marginal social benefits to different parties of market-based policies (i.e., some get more "clean air" than others). In the smaller, general resource allocation literature, distributional concerns arise from a related issue – the fact that high-value consumers of the good or service will purchase more, and low-value consumers less, if allocation occurs through prices. This is the advantage of a market; it is the source of welfare gains. But under certain conditions, like rationing during wartime, or in the aftermath of a natural disaster, willingness-to-pay may be an unjust allocation criterion.

Weitzman (1977) describes goods and services for which this is the case as a "class of commodities whose just distribution to those having the greatest *need*," (emphasis added) as distinguished from want or preference, "is viewed by society as a desirable end in itself." Within this formulation of the problem, the price system turns out to be the more effective way to allocate "needs" when preferences are heterogeneous or income distribution egalitarian; rationing is more effective under the opposite conditions.

Water for residential consumption is an example of a good, some portion of which may fall within Weitzman's need-based commodity regime (drinking, bathing, cooking), and some portion of which would not (swimming pools, large lawns). The standard approach to allocating water under conditions of scarcity, wisely, restricts these less necessary uses, not basic needs. Yet, the standard approach does not recognize heterogeneity in willingness-to-pay for these uses, and thus is likely to result in welfare losses when compared with management through prices.

2.2. Theoretical Gains from Drought Pricing

The current approach to drought management achieves a citywide required demand reduction uniformly restricting outdoor uses. The theoretical welfare gains from price-based municipal water regulation come from possible substitution within and across households. First, prices allow households to choose end-use consumption as they see fit, given their preferences (i.e., households could substitute some indoor for outdoor reductions). Thus, if indoor demand is anything but perfectly inelastic, the current approach creates a deadweight loss (DWL).

We describe this DWL in Figure 1, which maps stylistic linear demand curves for indoor and outdoor water use against a required demand reduction (on the horizontal axis).[3] The outdoor reduction mandated under the current approach (ΔQ_{reg}) creates a shadow price for outdoor consumption (λ) that is higher than the current marginal price of water (p_w). Under a market-clearing price (p_w^*), some of the citywide required reduction would take place indoors, and the shaded DWL from rationing water outdoors would disappear.

Additional welfare losses from the current approach come from disallowing substitution across households. Household willingness to pay for marginal units of water may be heterogeneous, even within the same type of use. Mutually beneficial trades could be made between households with strong and weak preferences for outdoor consumption. Figure 2 describes households with the same indoor demand curve, but different preferences with respect to outdoor demand. Here we assume that indoor demand is the least elastic portion of demand (C), and that for outdoor demand, there is a group of relatively elastic households (A), and a group of somewhat less elastic households (B). If households are heterogeneous, outdoor regulations not only drive a wedge between outdoor shadow prices and current marginal prices, but, since they are the same for all households, they also create variation in outdoor shadow prices across households. A market-clearing price would realize all potential gains from trade, eliminating the shaded DWL triangles. Thus, the DWL from rationing outdoor water has three components: some households consume too much outdoors, some consume too little outdoors, and everyone consumes too much indoors.

The implications for economic theory of rationing outdoor water during a drought are straightforward. Consider a market with two goods, water (w) and a composite consumption good (x). Let p_x=1. Water consumption has two components, indoor use (w_{in}) and outdoor use (w_{out}), and for all households (i), $w_{in_i} + w_{out_i} = w_i$. Households maximize utility, subject to a budget constraint and a constraint on outdoor water consumption imposed by a rationing policy during a summer drought in consumer i's market (1), where ϕ_i is the Lagrange multiplier on the budget constraint and ψ_i is the Lagrange multiplier on the outdoor water consumption constraint. The rationing policy we model is a uniform constraint on the number of days households may legally water outdoors each week. However, the constraint on the amount of outdoor water consumed ($\overline{w}_{out_i}$) varies by household, since conditional on watering, each household may use a different quantity of water.

$$\max_{\{w_{in_i}, w_{out_i}, x_i\}} \quad U_i(w_{in_i}, w_{out_i}, x_i) \qquad s.t. \quad \begin{array}{c} p_w w_i + x_i \leq y_i : \phi_i \\ w_{out_i} \leq \overline{w}_{out_i} : \psi_i \end{array}, \tag{1}$$

Assuming the typical assumptions for an interior solution, the first-order conditions from the maximization problem are given in (2) – (4):

$$\frac{\partial U_i}{\partial x_i} - \phi_i = 0 \tag{2}$$

$$\frac{\partial U_i}{\partial w_{in_i}} - p_w \phi_i = 0 \tag{3}$$

$$\frac{\partial U_i}{\partial w_{out_i}} - p_w \phi_i - \psi_i = 0 \tag{4}$$

Let $\lambda_i = p_w + \psi_i$, so that λ is the shadow price of outdoor water consumption (as in Figure 2). The utility maximization problem gives rise to an indirect utility function, $V_i(p_w, \lambda_i, y_i)$, which by Roy's identity gives the Marshallian demand functions $w_{in_i} = w_{in_i}(p_w, \lambda_i, y_i)$ and $w_{out_i} = w_{out_i}(p_w, \lambda_i, y_i)$.[4]

The DWL from the rationing of outdoor water is a function of the market clearing water price, p_w^*, needed to achieve each market's required aggregate demand reduction (5):

$$\sum_{i=1}^{N} w_{in_i}(p_w) + w_{out_i}(\lambda_i) = \sum_{i=1}^{N} w_{in_i}(p_w^*) + w_{out_i}(p_w^*), \tag{5}$$

[4] $w_{in_i} = \frac{\partial V_i}{\partial p_w} / \frac{\partial V_i}{\partial y_i}; w_{out_i} = \frac{\partial V_i}{\partial \lambda_i} / \frac{\partial V_i}{\partial y_i}$.

where N is the number of households in the market. We approximate the DWL as the loss in Marshallian consumer surplus from imposing $\bar{w}_{out_i}$, rather than charging p_w^*. In Figure 2, this is depicted as the sum of the DWL triangles for each household. We sum DWL over all consumers in a market (6). A technically correct estimate of DWL requires the calculation of compensating or equivalent variation. Our Marshallian consumer surplus estimates should be considered an approximation of DWL.[5]

$$DWL = \sum_{i=1}^{N} \{ (p_w^* - p_w) w_{in_i}(p_w) + (p_w^* - \lambda_i) w_{out_i}(\lambda_i) - \int_{p_w}^{p_w^*} w_{in_i}(z)dz - \int_{\lambda_i}^{p_w^*} w_{out_i}(z)dz \} \qquad (6)$$

2.3. Price and Non-price Water Regulation

To date, few studies have addressed municipal drought policies in this framework. Collinge (1994) proposes a theoretical water entitlement transfer system similar to the municipal-level price-based approach we analyze empirically. An experimental study simulates water consumption from a common pool and predicts that customer heterogeneity will generate welfare losses from command-and-control water conservation policies (Krause *et al.* 2003). Neither of these analyses estimates the magnitude of potential welfare gains, nor do they explore distributional implications in any depth. Renwick and Archibald (1998) estimate water demand elasticities by income quartile in Santa Barbara, California, and use these estimates to compare the distributional implications of price and non-price water conservation policies, but do not consider welfare impacts.

Two empirical studies have estimated the welfare losses from rationing outdoor water during a drought in a single city. The economic costs of a two-day-per-week sprinkling restriction in Perth, Australia are just under $100 per household per season (Brennan et al. 2007).

---

[5] In some of our empirical models, we use constant-elasticity demand functions. In these cases, the difference between our Marshallian consumer surplus estimate and Hicksian equivalent variation is likely to be small (West and Williams 2004).

Mandatory water restrictions in Sydney in 2004-2005 resulted in welfare losses of about $150 per household (Grafton and Ward 2008).[6]

3. Econometric Models

In part, the dearth of studies analyzing the welfare impacts of type-of-use restrictions in the residential sector is attributable to a variety of methodological challenges. Few data reliably disaggregate residential water consumption into its component uses. In addition, the presence of non-linear prices and the censoring of outdoor demand at zero complicate econometric analyses of this type. In our choice of econometric models, we draw on a well-developed literature on the price elasticity of residential water demand to meet these challenges.

Of the hundreds of water price elasticity studies published since 1960, a few are related closely to our current work.[7] . The first study to reliably separate household water consumption into its component end uses and estimate price elasticity for specific uses of water is Mayer *et al.* (1998). We use their data, but ask different questions and use different methods.

Many U.S. water utilities charge increasing-block prices to residential customers, creating piecewise linear budget constraints, under which marginal price and the quantity consumed are positively correlated. Structural models and instrumental variables (IV) models are two methods used to estimate demand under increasing-block prices, accounting for price endogeneity. Structural econometric techniques that treat piecewise-linear budget constraints in a manner consistent with utility theory derive from studies of the wage elasticity of labor supply

[6] In a related paper, Timmins (2003) compares a mandatory low-flow appliance regulation (a technology standard) with a modest water tax, using aggregate consumption data from 13 groundwater-dependent California cities. He finds the tax to be more cost-effective than the technology standard in reducing groundwater aquifer lift-height in the long run.

[7] Meta-analysis suggests that the central tendency of short-run elasticity estimates over the past four decades is about -0.3, and of long-run estimates about -0.6 (Dalhuisen *et al*. 2004).

[9] Kink point probabilities are, on average, 5 percent for two-block price structures, and they range from 1 to 3 percent, on average, for four-block price structures. We divide the kink probabilities evenly (for each household day) between the marginal prices on either side of the kink.

under progressive income taxation (Burtless and Hausman 1978), and have benefited greatly from the generalizations of Hanemann (1984) and Moffitt (1986, 1990). Five structural models of water demand under non-linear prices have been estimated: Hewitt and Hanemann (1995), Rietveld *et al*. (1997), Pint (1999), Olmstead (2007), and Olmstead *et al.* (2007).

For the estimation of end-use demand, in our case a pair of partial demand equations for indoor and outdoor consumption, the likelihood function in a structural approach would be complicated. The likelihood function for such models is constructed in part by using the block "cutoffs" (the quantities at which marginal price increases), based on total billing period water consumption, to determine the probabilities of consumption at all possible locations along the household's budget constraint (linear segments and kink points). The likelihood function, in our case, would include a total demand equation, which would determine marginal price, and separate end-use demand equations for indoor and outdoor consumption.

We develop an alternative approach here, due to the availability of water demand parameter estimates from an existing study using the same data (Olmstead *et al.* 2007). From the structural model, we derive the probability for each household of consuming at each possible marginal price on each day. Probabilities are calculated as functions of the structural parameter estimates, the data, and characteristics of each household's water price structure (number and magnitude of marginal and infra-marginal prices, as well as block cutoffs). We use these probabilities to estimate an expected marginal price, the sum of the products of marginal prices, times the probabilities of facing those prices.[9] Our price instruments are then calculated as the seasonal average, by household, of these daily probability-weighted prices. Appendix A describes the estimation of these price instruments in greater detail.

We begin with a test of the validity of our identification strategy by examining a model of total daily water demand, using probability-weighted marginal prices as instruments. Using the same data as Olmstead *et al.* (2007), we replicate their results for total water demand.[10]

The equation for this total demand function is (7), in which w_{total} is total daily water demand for household *i* on day *t*, p is the marginal water price for which we instrument, $\tilde{y}$ is virtual income, Z is a vector of daily and seasonal weather variables, and H is a vector of household characteristics.[11] The error structure comprises θ, a household heterogeneity parameter, and ν, the residual.

$$\ln w_{total_{it}} = \alpha \ln p_{it} + \mu \ln \tilde{y}_i + \delta Z_{it} + \beta H_i + \theta_i + \nu_{it}. \tag{7}$$

Or, in matrix notation:

$$\ln \mathbf{w}_{total} = \mathbf{X}'\boldsymbol{\gamma}_{total} + \mathbf{v}_{total} \tag{8}$$

Having tested the usefulness of the price instruments, we proceed with the estimation of end-use models. We adopt different models for indoor and outdoor demand, due to the fact that about half of outdoor demand observations are equal to zero (and we observe no such censoring of indoor demand). The indoor demand equation, identical to (8), except for the different dependent variable, is represented in (9).

$$\ln \mathbf{w}_{in} = \mathbf{X}'\boldsymbol{\gamma}_{in} + \mathbf{v}_{in} \tag{9}$$

For outdoor demand, censored at zero, we estimate a Tobit model, described in (10).

[10] Identification is possible because the estimated probabilities used to create the price instruments incorporate variation in price structures (number and magnitude of prices and block cutoffs), characteristics not incorporated in the second stage of our demand estimation.
[11] Virtual income is annual household income, plus the difference between total water expenditures if the household had purchased all units at the marginal price, and actual total expenditures. This standard technique treats the implicit "subsidy" of the infra-marginal prices as lump-sum income transfers, and is originally due to Hall (1973).

$$\mathbf{w}^*_{out} = \mathbf{X}'\boldsymbol{\gamma}_{out} + \mathbf{v}_{out}$$
$$w_{out_{it}} = \begin{cases} w^*_{out_{it}} \text{ if } w^*_{out_{it}} > 0 \\ 0 \text{ otherwise} \end{cases} \tag{10}$$

Dealing with endogenous prices in the Tobit framework involves one extra step to obtain unbiased estimates (Newey 1987). In the first stage, we estimate fitted prices as functions of the price instruments and all of the exogenous covariates. In the second stage, we include both the fitted prices and the residuals from the fitted price equation as independent variables.

4. Data

The data used in this study were collected by Mayer et al. (1998) for a study funded by the American Water Works Association Research Foundation. The data comprise 1,082 households in 11 urban areas in the United States and Canada, served by 16 water utilities. The water utilities in the sample are: Denver Water Department (CO); Eugene Water and Electric Board (OR); Seattle Public Utilities (WA) (and three other Seattle-area utilities: City of Bellevue Utilities, Northshore Utility District, Highline Water District); San Diego Water Department (CA); City of Lompoc (CA); Walnut Valley Water District (CA); Las Virgenes Municipal Water District (CA); Tampa Water Department (FL); City of Phoenix (AZ); Tempe Public Works Department (AZ); City of Scottsdale Water Department (AZ); and two utilities in the Regional Municipality of Waterloo Water Services Division (Waterloo and Cambridge) (ONT). No drought restrictions were in place in any of these cities during the data collection period.

The study conducted a one-time household survey, mailed to a random sample of 1,000 single-family residences in each city. (Cities served by multiple utilities still had only 1,000 households surveyed, divided among the multiple utilities). The survey achieved an average 45% response rate across cities, and in all cases but the city of Tempe, AZ, the random sample had water use characteristics that were not statistically different from the utility's population of single-family water use accounts. The researchers then drew a random sample of 125-150 households from among the returned surveys for end-use water demand data collection; they drew 100 households in each city, plus an additional 25-50 backup households in case some refused to participate. Out of the 1,100 households originally chosen, 40 refused to participate. Some were replaced with households randomly selected from the extra accounts, leaving a total sample of 1,084 households, from which we drop two that did not report some household characteristics used as independent variables in our analysis. Sampling procedures, response rates, and statistical tests for differences between respondents and single-family customers as a whole are described in further detail in Mayer et al. (1998).

The sample is appealing in that it is representative of U.S. urban and suburban single-family homes. The type of drought restriction we simulate, an outdoor watering restriction, is targeted at precisely this consumer group (as opposed to apartment buildings, for example), since this is where most outdoor consumption takes place. Due to the nature of the sample and the policy question we analyze, our estimates and conclusions would not be readily transferable to settings in which multi-family housing predominates (for example, drought pricing would not be a viable policy mechanisms where apartment-dwellers did not typically pay a water bill). The sample is also appealing in its broad geographic scope, across variable climates. However, these 16 water utilities represent a small fraction of U.S. residential water consumption, and results should be interpreted in this light.

Daily household water demand is observed over two periods of two weeks each, once in an arid season and once in a wet season. The data were collected between 1996 and 1998, but for each household, the two seasons of observation occurred within the same year. Daily demand data were gathered by automatic data loggers, attached to magnetic household water meters by utility staff and, thus, out of sight during water use. Total demand was disaggregated into its end uses using magnetic sensors attached to water meters. These sensors recorded water pulses through the meter, converting flow data into a flow trace, which detects the "flow signatures" of individual residential appliances and fixtures (Mayer *et al.* 1998). We add together consumption from all indoor fixtures (primarily toilets, clothes washers, showers, and faucets) to obtain indoor demand, and consumption from all outdoor uses (irrigation and pools) to obtain outdoor demand.[13] Leaks and unknown uses are included in total demand, but are not modeled explicitly as either indoor or outdoor consumption.[14]

We use the Mayer et al. (1998) survey data for household characteristics, including gross annual household income, family size, home age and size, lot size, the number of bathrooms, and the presence of evaporative cooling.[15] Daily weather observations are from local weather stations. Control variables are: season (arid vs. wet), maximum daily temperature, and evapotranspiration less effective rainfall (0.6 times total rainfall). Finally, regional fixed effects control for climate variation not absorbed by daily and seasonal weather variables.[16]

---

[13] The listed indoor uses comprise 94 percent of indoor consumption, on average. Remaining indoor uses include evaporative cooling, dishwashers, bathtubs, water treatment, hot tubs, and humidifiers. In outdoor use, we can distinguish between water consumption for swimming pools from that for all other outdoor uses. We cannot distinguish among irrigation, car-washing, and washing of sidewalks and driveways, but these uses are all typically regulated or prohibited by drought policies.

[14] Leaks comprise approximately 6 percent of total consumption, and unknown uses approximately 1 percent.

[15] Evaporative cooling, common in arid climates, substitutes water for electricity in the provision of air conditioning. Less than 10 percent of sample households have evaporative coolers, but 43 percent of sample households in Phoenix have them, and about one-third of households in Tempe and Scottsdale. Households with evaporative cooling use, on average, 35 percent more water than households without.

[16] Regions are: (1) Southern California (Las Virgenes MWD, City of San Diego, Walnut Valley Water District, City of Lompoc); (2) Arizona and Colorado (Phoenix, Tempe, Scottsdale, Denver); (3) Northern (City of Seattle Public Utilities, Highline Water District, City of Bellevue Utilities, Northshore Utility District, Eugene Water and Electric Board, Regional Municipality of Waterloo, Ontario); (4) Florida (City of Tampa Utilities).

Table 1 provides descriptive statistics. Water demand varies by season, but only for outdoor use. Outdoor water demand in an arid season is, on average, five times outdoor demand during a wet season. In addition, the fraction of observations using any water outdoors, at all, is 0.42 – the reason for choosing a censored regression model. Note that while we observe zero indoor water consumption on rare occasions (10 household-days), more than 99% of household-days have some consumption indoors.

The households in the sample face either uniform marginal prices (39 percent); or two-tier (44 percent) or four-tier (17 percent) increasing block prices.[17] Each household faces one price structure throughout each season of observation, but six sample utilities changed prices or price structures between the two periods.[18] Given cross-sectional and time series variation, there are 26 price structures in the data; eight two-tier increasing block structures, ten four-tier increasing block structures, and eight uniform marginal prices. Price variation in the sample is primarily in the cross-section. Regressing prices on city fixed effects results in an R-squared of 0.75, on household fixed effects, 0.94, and on our price instruments, 0.88.

Marginal prices range from \$0.00 per thousand gallons (kgal) for the first 4,490 gallons per month in Phoenix, to \$4.96 per kgal in the most expensive block in the Las Virgenes Municipal Water District, with an average marginal price of \$1.71/kgal. The mean of our price instrument, $\widehat{p}$, is equal to the mean of observed marginal water prices, and the standard

[17] About one-third of households in the United States face increasing-block prices. Thus, these price structures are over-sampled in the data. This matters for elasticity estimates only if elasticity varies across price structures – an unresolved empirical question (Olmstead *et al.* 2007).
[18] The utilities with price changes across the two seasons were: San Diego, Scottsdale, and all four utilities in the Seattle area (City of Seattle, Bellevue, Highline, and Northshore).

deviation is slightly smaller.[19] Average total expenditures on water in the sample, including fixed charges, are \$326 per year, or about 0.47 percent of average annual household income.

In our tests of consumer heterogeneity, we divide the sample into four subgroups, based on income and lot size. Income is our best available proxy for ability to pay, and lot size is our best available proxy for preferences for the services that households derive from outdoor water consumption (such as lawns, gardens, pools, and looking better than the neighbors). We expect that wealthier consumers and those with larger lots will be less price sensitive. Those with both incomes and lot sizes above the medians (\$55,000 per year, and 9,000 ft^2, respectively) are categorized as "rich, big lot" households; those with both incomes and lot sizes below the medians are categorized as "poor, small lot" households; and so on for the two groups in between. In the absence of any drought policy these groups divide total sample water consumption as follows: rich, big lot (43 percent); rich, small lot (23 percent); poor, big lot (15 percent); poor, small lot (19 percent). Households may also be heterogeneous within groups; in this sense, we will underestimate the true DWL from rationing.

5. Results

[19] For some sample utilities, marginal wastewater charges are assessed on current water consumption. In addition, some sample utilities benchmark water use during the wet season as a basis for volumetric wastewater charges assessed the following year. For households observed during these periods (and there are some in the data), effective marginal water prices would include some function of the present value of expected future wastewater charges associated with current use. We do not do this here; marginal wastewater charges are excluded from the present analysis.

[24] The Tobit price coefficients are not elasticities – see the notes to Table 4 for the calculation of Tobit elasticities. For these estimates, we calculate the Tobit model allowing different coefficients on price for the arid and wet seasons.

5.1. End-Use Price Elasticity

We begin by estimating a total water demand model as in (8), using our 2SLS approach. Table 2 reports coefficient estimates and standard errors from two such models. The first column contains estimates from Olmstead *et al.* (2007), for the purpose of comparison. The second column reports estimates from a 2SLS random-effects model for panel data, in which the independent variables in the demand function are identical to those in Olmstead *et al.* (2007). This model generates parameter estimates that are similar to those from the structural model, with an important exception – the price elasticity estimate is not significantly different from zero.

In the third column of Table 2, we group the city-level fixed effects from “test model 1” into regional fixed effects. This model captures exogenous sources of geographic and climatic variation, leaving enough price variation to identify a price elasticity. With regional fixed effects, we obtain estimates that are very close to those from the structural model, with less precision. The price elasticity estimate is -0.36, and strongly significant. We use test model 2, a 2SLS random-effects model with regional fixed effects, as our point of departure for the rest of the analysis.

We separate total demand into indoor and outdoor consumption and estimate the models given in (9) and (10). Table 3 reports the full set of parameter estimates and standard errors (with the exception of regional fixed effects) for indoor and outdoor demand models, both annually and by season.

Indoor use appears to be influenced by income, family size, and evaporative cooling. Outdoor demand parameters are all significant, with the exception of home age. Many significant outdoor demand parameters would seem to be drivers of indoor, rather than outdoor consumption (for example, the number of bathrooms). It may be that these variables are correlated with omitted variables that represent preferences for outdoor water consumption.

Table 4 summarizes the results of these models with respect to price, the parameter of interest, and reports elasticities, rather than price coefficient estimates, for the outdoor models.[24] The partial demand models reveal striking variation in elasticity across uses and, for outdoor use, across seasons. None of the indoor elasticity estimates are significantly different from zero (the estimates are very small in magnitude, as well). Outdoor demand in the wet season is the most price elastic (-1.18), and is still quite responsive to price in the arid season (-0.74). Essentially all of the strong seasonal variation in total water consumption is attributable to outdoor use.[25]

Does targeting outdoor consumption, as command-and-control water conservation policies do, provide a good first approximation to a price-based approach? Indeed, outdoor uses are the uses that households, themselves, would choose to cut back the most in response to a price increase. But an important cost of the prescriptive approach derives from household heterogeneity.

5.2. Household Heterogeneity

To test whether households are, in fact, heterogeneous in their preferences for water consumption, we divide the sample into four sub-groups, based on income and lot size.[26] We estimate separate elasticities for the four groups. Results, reported in Table 5, indicate a high degree of heterogeneity. For all customer groups, the price elasticity of indoor demand is insignificantly different from zero. Households presumed to have the strongest preferences for outdoor water consumption, the "rich, big lot" group, exhibit the least elastic outdoor demand (-

[25] We have no information on available substitutes for municipal tap water in outdoor uses, which might include groundwater wells or public surface water sources. To the extent that these substitutes are available in the sample, our outdoor elasticity estimates are greater in magnitude than they would be in the absence of substitutes.
[26] We define these heterogeneous groups based on sample median income and lot size, and test an alternative definition of city-specific median lot size as a robustness check (see Appendix B).

0.48). Those presumed to have the weakest preferences for outdoor consumption, the "poor, small lot" group, exhibit the most elastic outdoor demand (-0.87). The two middle groups appear to be about equally price elastic.[27]

5.3. Robustness of Elasticity Estimates

We test the robustness of our elasticity estimates by exploring a number of other model specifications. These include demand functions with household fixed effects and functions with city fixed effects (rather than regional fixed effects). We also examine a model that collapses the daily variation in household demand, obtaining parameter estimates from regressions of aggregate seasonal household demand on the covariates of interest. Our results are qualitatively robust to these alternative models and others. See Appendix B for details.

6. Simulations and Discussion

Traditional regulations typically limit the number of days in a week that households may use water outdoors (watering lawns, washing cars, or filling swimming pools). The stringency and enforcement of these type-of-use restrictions vary greatly (Dixon *et al*. 1996). A common policy is to limit outdoor watering to two days a week. We examine the implications of this policy, as well as limits of three, one, and zero days per week.

Households' willingness to pay for the marginal unit of water should increase with drought policy stringency. To calculate shadow prices, we estimate the constrained level of expected consumption for each household under each policy, and then "back up" along that household's demand curve – using equation (10) – to obtain their willingness-to-pay for the marginal unit of water.[28] Some households are unconstrained by the policies; their probability of watering on a given day is less than or equal to the probability imposed by the watering

---

[27] F-tests find "rich, big lot" to differ significantly from "rich small," "poor big," and "poor small" (p-values are 0.01, 0.02, and 0.01, respectively). The other categories do not differ significantly from each other.
[28] We estimate shadow prices using the outdoor price coefficient underlying the elasticity in Table 5.

restrictions. For constrained households, we calculate the difference in their expected quantity demanded in the unrestricted and restricted scenarios.

For example, for a twice-per-week watering policy, the restricted probability of watering is 2/7 (assuming full compliance). A household with a probability of watering greater than 2/7 is constrained, and will have a resulting decrease in expected quantity demanded, $E(\Delta w_{out_{it}})$, as in (11). Let:

$$\begin{aligned}\mathbf{prob} &= \Pr(\mathbf{X}'\hat{\boldsymbol{\gamma}}_{out} + \mathbf{v}_{out} > 0) \\ \mathbf{w}_{out}^{cond} &= E(\mathbf{X}'\hat{\boldsymbol{\gamma}}_{out} + \mathbf{v}_{out} \mid \mathbf{X}'\hat{\boldsymbol{\gamma}}_{out} + \mathbf{v}_{out} > 0)\end{aligned}$$

Then for each household day:

$$E(\Delta w_{out_{it}}) = \begin{cases} (prob_{it} - \frac{2}{7}) * w_{out_{it}}^{cond} \text{ if } prob_{it} > \frac{2}{7} \\ 0 \text{ otherwise} \end{cases} \quad (11)$$

We estimate $E(\Delta w_{out_{it}})$ for the arid season only, for the four drought policies described above.

We assume full compliance with the simulated drought policies, since there are no watering restrictions in effect our data, and no estimates are available from the literature on the typical water savings achieved with such policies. Adding a probability of compliance that is less than one to the analysis would be equivalent to simulating different drought policies than those we select.[29] It is likely that compliance will be higher under a price-based policy than under the current prescriptive approach.[30] Nevertheless, were utilities to estimate market-clearing prices,

[29] If non-compliance is punished with a fine, some of the benefits of drought pricing may be realized under the rationing policy, depending on the magnitude of the fine.

[30] "Cheating" in the market context would require that households figure out how to consume piped water outdoors off-meter, and in the current context can easily be accomplished by watering at night, or in some other way that avoids observation by utility staff or vengeful neighbors.

they might want to charge slightly higher prices, anticipating less than full compliance, depending on the importance of meeting the demand reduction.[31]

Second, we assume conditional outdoor demand is unchanged under the simulated drought policies – households water less frequently, in compliance with the policy, but the same amount per watering as they did in the absence of regulation. Most municipal drought ordinances forbid allowing water to run off of residential properties onto sidewalks and streets, making it unlikely that households would over-water in response to reduced allowable watering frequency. Nonetheless, if conditional demand were to increase under the drought policies we simulate, it would decrease the aggregate demand reduction achieved. Thus, like full compliance, our assumption regarding conditional demand could be altered simply by simulating different drought policies than those we choose. Were we to deal with both of these issues by simulating a set of percentage aggregate demand reductions, rather than allowable frequency of watering, we would lose an important benefit of our current approach. The apportionment of aggregate reductions across households would be arbitrary, whereas now it is based upon a household's probability of watering, calculated using the data and parameter estimates.

Finally, note that we do not simulate the impact on demand and welfare of an actual drought – we make no changes to evapotranspiration, rainfall, or maximum daily temperature. Simulating an actual drought would increase the welfare impact of moving to a price-based approach, due to the reduced availability of a substitute (rain), but not by much. Even a 25 percent increase in our weather variable that describes outdoor watering needs (*weath*) would

[31] This is always the case with a market approach based on prices, rather than quantities (Weitzman 1973). While a quantity instrument (like tradable permits) would be preferable in cases where utilities had very strict quantity constraints, such as the threat of violation of a treaty over shared water resources, the transaction costs involved in establishing a household-level trading regime would likely be prohibitive.

increase consumption by less than one percent.[32] In addition, the characteristics of a drought vary significantly across sample cities, depending not on weather variables, alone, but also reservoir capacity and other characteristics. Simulating drought, and not just drought policy, in this wide cross-section of North American cities would be somewhat arbitrary.

6.1. Shadow Prices

The evidence of heterogeneity among households from Section 5.2 suggests that, when outdoor water consumption is restricted by drought policy, shadow prices for the marginal unit of water will vary significantly. Variation in shadow prices would indicate potential gains from trade. Based on the separate elasticity estimates for our household sub-groups, we calculate shadow prices (by household-day) and market-clearing prices (by utility) under drought policies of varying stringency.

Table 6 reports shadow prices for the arid season. In the most extreme policy, when no watering is allowed, our nonlinear functional form implies an infinite shadow price for all customers. We assume that the willingness-to-pay is at most $50 per thousand gallons. The most common policy (of allowing outdoor watering two days per week) has an average shadow price of $4.98 per thousand gallons. Note that this is almost three times the average marginal price consumers actually pay ($1.71).

As we would expect, shadow prices increase with the stringency of the drought policy. Furthermore, the standard deviation of shadow prices across all customers is increasing in drought policy stringency. These standard deviations reflect the potential benefits from establishing a price-based policy. For example, Figure 3 provides histograms of shadow prices in

[32] Our calculation is based on the sample average and coefficient estimate for *weath* reported in Tables 1 and 2, respectively. Note the substantial variation in *weath*; a 25 percent increase is not out-of-sample.

two cities given a policy limiting outdoor watering to two days a week.[33] Even in cities with relatively small standard deviations in shadow prices (like Eugene, Oregon), there is a lot of variation, and shadow prices tend to be right-skewed. There are some households with much higher shadow prices than the average. Under drought pricing, all households will consume such that shadow prices are equal.

6.2. Market-clearing Prices

We then construct market-clearing prices under drought policies of varying stringency. Here we assume that each utility's goal is simply to save the aggregate quantity of water it would save by implementing each type of drought policy, no matter how that aggregate water consumption reduction is achieved. The aggregate water savings implied by each drought policy is the sum over the $E(\Delta w_{out_{it}})$ in (11), for all households within a utility. From the households' perspective, these savings could be achieved indoors, outdoors, or by purchasing "credits" from another household.[34, 35] We then identify the market-clearing price for that aggregate reduction, constraining households to non-negative consumption.[36]

In Figure 3, the solid vertical lines denote the market-clearing prices for two cities. Differences in average shadow and market prices are created by the skewness of shadow prices and magnified by their variation. In Eugene, where there was relatively little shadow price

---

[33] We chose two cities (Eugene, Oregon, and San Diego, California) that provide examples with low and average amounts of variation in shadow prices, respectively.

[34] An actual tradable credit system would likely be infeasible due to large transactions costs. However, with no uncertainty, a regulator could equivalently set a higher price so as to clear the market.

[35] Price coefficients for indoor and outdoor demand underlying the indoor and outdoor elasticities reported in Table 5 are used for estimation of shadow and market prices. While indoor elasticity estimates in our models are not significantly different from zero, not all indoor uses are perfectly inelastic. In separate demand equations for indoor uses, we find that showers and clothes-washing each have elasticities of about –0.22 (significant at .05), using Tobit models for both due to censoring. The other indoor uses have demand curves that are approximately vertical. Given evidence of some price-responsiveness indoors, and the small likelihood that utilities would or could regulate specific indoor uses, we use our overall indoor elasticity estimates in the welfare analysis.

[36] Market-clearing prices are estimated for the irrigation season, assuming that drought regulations are implemented primarily during arid months.

variation, the market-clearing price is close to the average of the shadow prices. In San Diego, we see a larger difference between the shadow prices and the market price.

The last column of Table 6 reports market-clearing prices by utility. Like the shadow prices, market-clearing prices increase monotonically with the stringency of watering restrictions. Within a utility, there is a common price. Market-clearing prices vary substantially, though not as much as shadow prices, across utilities. For the most common drought policy, the average market-clearing price is \$3.90 per thousand gallons, or slightly more than twice the current mean marginal price.

6.3. Welfare Implications

The management of water scarcity through residential outdoor watering restrictions results in substantial welfare losses, given the observed heterogeneity in willingness-to-pay. Welfare losses calculated in this context (with a reduced-form model, and Marshallian demand curves) should be considered very rough estimates. Nonetheless, we do calculate them.[37]

For each utility, we simulate the welfare effects of a two-day per week watering policy over a 180-day irrigation season. Deadweight losses (DWL) under the current regime represent the estimated benefit to the average household of introducing drought pricing.[38] Table 7 reports the median estimate of average per-household DWL by utility, with 5th and 95th percentiles in parentheses, from 1500 replications of a nonparametric bootstrap. For the bootstrap sampling, we cluster by household, utility and season to account for the fact that observations for a household across seasons are not independent, and to preserve the price variation across utilities and

[37] We estimate deadweight loss by integrating demand curves, as discussed in Table 7.

[38] Presumably, if we included multi-family homes in a market-based policy, gains from trade would be somewhat larger. Their inclusion would add their (currently unregulated) indoor use to the "common pool" which, if sensitive to price increases, would be an additional source of reductions to support increases in higher-valued uses.

[40] Without bootstrapping, the average per-household DWL for the sample is about \$98 per household; a small upward bias would be introduced from attributing some variation in shadow prices that is actually sampling error to household heterogeneity.

seasons from the original sample. We re-estimate indoor and outdoor demand elasticities by household group, as well as shadow prices, market-clearing prices, and DWL by utility. The median of bootstrapped DWL estimates ranges from $2.30 per household in the service area of Seattle Public Utilities, to $407.50 per household in the service area of the Las Virgenes Municipal Water District. The variation in DWL from the current regulations is attributable, in part, to the standard deviation of shadow prices, a strong indicator of potential gains from trade. In our sample, society would be better off by an average of $92 per household through the introduction of drought pricing.[40] About 61 percent of this DWL, or $56 [37, 89], comes from allowing households to substitute indoor for outdoor reductions, and the remainder from allowing substitution across households. We have defined heterogenous household preferences in terms of four groups of households; in reality, preferences and elasticities likely vary by household, though our data are insufficient to estimate elasticities at this fine level. Our DWL estimates would increase with a more complete picture of heterogeneity in the sample.[41]

The discussion of price-based approaches is largely hypothetical under the current regulatory structure. Estimated market-clearing prices are greater than current average marginal prices in all of these markets, in some cases by a lot. If they are also higher than average costs, prices this high would be impossible to implement without significant rebates of some form, as we discuss in Section 6.5.

In addition, we face the standard worry about the welfare effects of a theoretical first-best policy in a second-best setting (Lipsey and Lancaster 1956, Harberger 1974). Ours is a partial equilibrium analysis, thus spillovers (no pun intended) to other markets from changes in water expenditures as a result of a policy change may be a concern, as it is in the well-known

[41] For example, if we divide households by treciles of income and lot size, creating nine heterogeneous groups, we estimate an average per-household DWL (and 90% confidence interval) across the sample of $118 [71, 218].

theoretical and empirical studies of environmental taxation in a second-best setting (Sandmo 1975, Goulder *et al.* 1999, Goulder and Williams 2003).

In our case, the distortions of greatest concern are within water markets, themselves. In most cases, marginal water prices are well below the marginal social cost of water supply (Hanemann 1997, Timmins 2003). Drought pricing will result in higher marginal prices for all households (even if total expenditures fall for some households through lump-sum transfers). To the extent that drought pricing results in more households paying something closer to marginal social cost, the pre-existing distortion does not change the basic nature of our results.[42] However, if marginal prices in the sample are well below marginal social cost, the welfare impacts of moving to drought pricing may pale in comparison to the impacts of moving to marginal social cost pricing, period. This is an important issue, but it is beyond the scope of this analysis.

6.4. Distributional Implications

While the shift from outdoor watering restrictions to a price-based municipal drought policy would be welfare-improving in all markets, the distributional implications depend on the allocation of property rights. Prices would re-distribute scarce water so that those with high willingness-to-pay for water consumption outdoors would consume more than they do under outdoor use restrictions, and those with low willingness-to-pay would consume less. Hence, drought pricing would result in a water allocation that would "soak the rich."

Under drought pricing, relative to the traditional approach, the consumption share of the least elastic group (at least for outdoor uses), the rich, big lot households, would rise from 35 to 48 percent; the consumption share for the most elastic group, the poor, small lot households, would fall from 23 to 16 percent, with smaller reductions in consumption shares by the

[42] An increase in the marginal price of municipal water supply will generate an increase in the consumption of substitutes. Where groundwater is a viable substitute for municipal tap water, spatial and intertemporal externalities may result (or increase, where they are already present).

remaining two groups. Absolute consumption falls quite drastically among all groups under both types of drought policies.

Given these changes in consumption shares, the largest DWL under the current approach is experienced by the rich, big lot households (Table 8). But a more meaningful number to households is the change in surplus. We calculate the average changes in producers' surplus (PS) and consumers' surplus (CS), as well as the average change in CS by group that would result from the adoption of a price-based approach. On average, consumers in each group are worse off under drought pricing, primarily because current average marginal prices do not reflect scarcity, and are thus below current shadow prices. Thus, consumers would not support this efficiency-improving change without a rebate. The average PS, \$136 per household, could be used for this purpose, as utilities in the United States are usually restricted to zero (or very small) profits.[43] Households' minimum rebate to support the policy may equal their average change in CS. The ordering of DWL by household group follows the ordering of consumption shares; those groups who use the most water experience the greatest losses from rationing. (The ordering of changes in CS is similar, but flips relative to consumption shares for the last two groups.)

Conditional on lot size, drought pricing without a rebate would be regressive (column 3 in Table 8). A progressive price-based approach can always be designed through the use of transfers. In the present case, this could occur through the utility billing process. Should utilities seek to address distributional concerns through the water supply system, rebates could be delivered based on income. Low-income households could receive sums greater than their minimum willingness-to-accept, since the average PS exceeds the average CS by the size of the average re-captured DWL. Market-based policies need not be regressive.

[43] Note that the sum of changes in average CS and PS equal the average DWL, by definition.

7. Conclusions

Using unique panel data on residential end-uses of water, we examine the welfare implications of outdoor water rationing as a demand reduction policy. Using price variation across and within markets, we identify price elasticities for indoor and outdoor consumption. Outdoor uses are more elastic than indoor uses, suggesting that current policies target those water uses households, themselves, are most willing to forgo. Nevertheless, we find that use restrictions have substantial welfare implications, primarily due to household heterogeneity in willingness-to-pay for scarce water.

Heterogeneity is often ignored in economic analyses, which proceed from the viewpoint of the "representative consumer." For heavily regulated goods, estimating the welfare gains from introducing markets requires the opposite starting point—it is precisely the variation in marginal benefits that opens up potential gains from trade within non-market allocations. We find some potential for substitution within households across end-uses, and some from substitution across households.

Of all the currently regulated markets in which alternative price-based policies have been proposed, municipal water markets may be the easiest in which to imagine actually introducing a price-based approach, even one that involves lump sum transfers to achieve equity goals. Household water use is metered, and monitored by utility staff for the purpose of billing and collection.[44] Were such a system to be implemented, a municipality would have the rare opportunity to affect an actual Pareto improvement, in which gains not only exceed losses, but no household is made worse off.

[44] This is quite unlike the case of market-based pollution regulation, which requires the installation of continuous emissions monitoring infrastructure (for tradable permits), or the case of congestion pricing, which requires a new system with which regulators can track consumers' use of priced roadways.

If concern about "everyone doing their part" during a drought is the reason for the current predominance of command-and-control, rather than market-based approaches to the management of scarce water resources, economists' discussion of potential lump sum transfers and actual Pareto improvements may fall on deaf ears. There is irony in this. In the long run, command-and-control regulations provide no incentive for the invention, innovation, and diffusion of water conserving technologies (outdoors or indoors). Water priced below marginal social cost also results in inefficient land-use patterns, like the establishment of large, lawn-covered lots and thirsty non-native plant species where water is scarce. Further investigation of the welfare gains from water marketing within and across sectors is an important area for further research.

Appendix A. Estimation of Price Instruments

The water demand function (A.1) is in exponential form, where w is total daily water demand, Z is a matrix of seasonal and daily weather conditions, X is a matrix of household characteristics, p is the marginal water price, $\tilde{y}$ is virtual income, η is a measure of household heterogeneity, ε is optimization or perception error; and δ, β, α, and μ are parameters.[45]

$$w = e^{Z\delta} e^{X\beta} p^{\alpha} \tilde{y}^{\mu} e^{\eta} e^{\varepsilon} \qquad \text{(A.1)}$$

Let $\underline{w}_k^*(.) = e^{Z\delta} e^{X\beta} p_k^{\alpha} \tilde{y}_k^{\mu}$, or optimal consumption in block k. Then, unconditional demand under a two-tier increasing-block price structure, in which w_1 is the kink point, can be represented as in (A.2), and unconditional price as in (A.3).

[45] The structural model includes two additional parameters, σ_η and σ_ε. Our 2SLS approach does not allow separate identification of the two error variances. We use the structural estimate of σ_η to calculate block and kink probabilities in (A.5).

$$w = \begin{cases} \underline{w}_1^*(.)e^{\eta}e^{\varepsilon} & if \quad 0 < e^{\eta} \le \dfrac{w_1}{\underline{w}_1^*(.)} \\ w_1 e^{\varepsilon} & if \quad \dfrac{w_1}{\underline{w}_1^*(.)} < e^{\eta} \le \dfrac{w_1}{\underline{w}_2^*(.)} \\ \underline{w}_2^*(.)e^{\eta}e^{\varepsilon} & if \quad \dfrac{w_1}{\underline{w}_2^*(.)} < e^{\eta} \end{cases} \tag{A.2}$$

$$p = \begin{cases} p_1 & if \quad 0 < e^{\eta} \le \dfrac{w_1}{\underline{w}_1^*(.)} \\ \text{indet.} & if \quad \dfrac{w_1}{\underline{w}_1^*(.)} < e^{\eta} \le \dfrac{w_1}{\underline{w}_2^*(.)} \\ p_2 & if \quad \dfrac{w_1}{\underline{w}_2^*(.)} < e^{\eta} \end{cases} \tag{A.3}$$

Consumption only occurs at the kink point if the consumer maximizes utility for choices that are unavailable at all (p_k,y_k), so for kink observations, $\underline{w}_1^*(.) > w_1$ and $\underline{w}_2^*(.) < w_1$.

From the conditional price equation, we derive a daily probability-weighted price (A.4). Our price instrument is the seasonal average, by household, of $\widehat{p}$. Errors are assumed to be independent and normally distributed. Thus, $e^{\eta} \sim LN(\mu_{e^{\eta}}, \sigma_{e^{\eta}})$, and integrations in (A.5) are over the probability density function of this lognormal distribution.

$$\widehat{p} = \Pr A * p_1 + \Pr B * (.5 p_1 + .5 p_2) + \Pr C * p_2 \tag{A.4}$$

Where:

$$\begin{aligned} &\Pr A = \int_0^{\frac{w_1}{\underline{w}_1^*(.)}} f(e^{\eta})de^{\eta} \\ &\Pr C = \int_{\frac{w_1}{\underline{w}_2^*(.)}}^{\infty} f(e^{\eta})de^{\eta} \\ &and \quad \Pr B = 1 - \Pr A - \Pr C \end{aligned} \tag{A.5}$$

Appendix B. Robustness of Estimation

The ideal data for this analysis would include a longer time-series component. Essentially, we have two price observations per household. For this reason, we cannot estimate a model with household fixed effects (FEs), or even city FEs, without reducing price variation so substantially as to prevent reasonable interpretation of parameter estimates, provided we are able to estimate effects, at all, that are significantly different from zero.

The best way we can control for household heterogeneity in this context without losing too many degrees of freedom is to include a household random effect in the models. Hausman tests in most cases reject the null hypothesis that the random effects model is consistent and efficient. However, where we reject random effects, we do so largely because the estimates are more precise than they are in models in which the Hausman test does not reject the null.

As tests of robustness, we estimate models with household and city FEs, as well as those using aggregate data. For the latter, we collapse daily observations to seasonal observations, creating two demand observations per household, and obtain estimates from regressions of aggregate seasonal household demand on the covariates.

The household fixed effects (FE) model generates an upward-sloping indoor demand function, though the price coefficient is insignificant; this is likely an artifact of the small amount of price variation available to estimate the indoor elasticity. This is true both for the homogeneous demand functions, and for those allowing household heterogeneity. The incidental parameters problem prevents us from estimating maximum likelihood outdoor models with household FEs.

Table B.1. reports results of the city FE and aggregate demand models for homogeneous households. With city FEs, indoor demand is upward-sloping (though the elasticity is also insignificant), and the magnitude of the outdoor elasticity is much smaller than the magnitude of the outdoor elasticity with regional FEs. For this sample, the inclusion of city FEs substantially reduces price variation and results in noisy estimates of price elasticities. Fixed effects of an even finer level, that of utilities, might control for residential water conservation programs—and other utility policies—unlike either the regional or city FEs. The inclusion of regional fixed effects does control for long-run climate variation but cannot capture utility-specific heterogeneity.

The aggregate demand models of indoor and outdoor consumption provide estimates that are very similar to the daily models. That our models are robust to this seasonal specification is encouraging. However, the daily water demand models are preferable in that they provide more detailed information regarding the impact of daily weather conditions on water consumption.

We run this same set of robustness checks (household FEs, city FEs, and aggregate data) on our estimates of elasticity across heterogeneous consumer groups, reported in Table B.2. Surprisingly, the city FE models are able to identify outdoor demand elasticities for each consumer group – they are approximately equal, in statistical terms. Results by group using aggregate data differ slightly from our daily models. Using aggregate data, the difference between the "rich, big lot" households and the other groups is more pronounced. The estimates for the remaining three groups are not significantly different from each other.

Finally, we also test an alternative definition of household heterogeneity. The logic of the sample-level definition of heterogeneity is that wealthy households in one city are more like households of similar income elsewhere than those of lower income within a city, and that lot size primarily reflects preferences for outdoor water-intensive activities and services. However, differences in land values across cities suggest that an alternative model would consider lot size heterogeneity relative to city medians, rather than the sample median. We estimate elasticities for household groups using city median lot size and sample median income, and report results in the last column of Table B.2. Elasticities are somewhat less different across groups under this specification. The average DWL in this case is about $34 per household. In general, we find the demand models, in aggregate and by group, to be robust to the alternative specifications tested.

References

Baumol, William J. and Wallace E. Oates. 1988. *The Theory of Environmental Policy, Second Edition*. New York: Cambridge University Press.

Brennan, D., S. Tapsuwan, and G. Ingram. 2007. The welfare costs of urban outdoor water restrictions. *Australian Journal of Agricultural and Resource Economics* 51: 243-261.

Burtless, Gary and Jerry A. Hausman. 1978. The Effect of Taxation on Labor Supply: Evaluating the Gary Income Maintenance Experiment. *Journal of Political Economy* 86: 1101-1130.

Burtraw, Dallas, Alan Krupnick, Erin Mansur, David Austin, and Deirdre Farrell. 1998. The Costs and Benefits of Reducing Air Pollution Related to Acid Rain. *Contemporary Economic Policy* 16: 379-400.

Collinge, Robert A. 1994. Transferable Rate Entitlements: The Overlooked Opportunity in Municipal Water Pricing. *Public Finance Quarterly* 22: 46-64.

Crocker, Thomas D. 1966. "The Structuring of Atmospheric Pollution Control Systems." In *The Economics of Air Pollution*, ed. H. Wolozin, 61-86. New York: Norton.

Dales, John H. 1968. *Pollution, Property and Prices*. Toronto: University of Toronto Press.

Dalhuisen, Jasper M., Raymond J. G. M. Florax, Henri L. F. de Groot, and Peter Nijkamp. 2003. Price and Income Elasticities of Residential Water Demand: A Meta-analysis. *Land Economics* 79: 292-308.

Daniel, Joseph I. 2001. Distributional Consequences of Airport Congestion Pricing. *Journal of Urban Economics* 50: 230-258.

De Vany, Arthur S. and W. David Walls. 1999. Price Dynamics in a Network of Decentralized Power Markets. *Journal of Regulatory Economics* 15: 123-40.

Dixon, Lloyd S., Nancy Y. Moore, and Ellen M. Pint. 1996. *Drought Management Policies and Economic Effects in Urban Areas of California, 1987-1992*. Santa Monica: RAND.

Gilbert, Richard, Edward Kahn, and Matt White. 1996. "The Efficiency of Market Coordination: Evidence from Wholesale Electric Power Pools." In *Networks, infrastructure, and the new task for regulation*, eds. W. Sichel and D. L. Alexander, 37-58. Ann Arbor: University of Michigan Press.

Goulder, Lawrence H., Ian W. H. Parry, Roberton C. Williams III, and Dallas Burtraw. 1999. The Cost-Effectiveness of Alternative Instruments for Environmental Protection in a Second-best Setting. *Journal of Public Economics* 72: 329-360.

Goulder, Lawrence H. and Roberton C. Williams III. 2003. The Substantial Bias from Ignoring General Equilibrium Effects in Estimating Excess Burden, and a Practical Solution. *Journal of Political Economy* 111: 898-927.

Grafton, R. Quentin and Michael Ward. 2008. Prices versus rationing: Marshallian surplus and mandatory water restrictions. *Economic Record* 84: S57-S65.

Hall, Robert E. 1973. "Wages, Income and Hours of Work in the U.S. Labor Force." In *Income Maintenance and Labor Supply*, ed. G. G. Cain and H. W. Watts, 102-159. Chicago: Rand McNally College Publishing Company.

Hanemann, W. Michael. 1984. Discrete/Continuous Models of Consumer Demand. *Econometrica* 52: 541-561.

Hanemann, W. Michael. 1997. "Price and Rate Structures." In *Urban Water Demand Management and Planning*, ed. D. D. Baumann, J. J. Boland, and W. M. Hanemann, 137-179. New York: McGraw-Hill.

Harberger, Arnold C. 1974. *Taxation and Welfare*. Chicago: University of Chicago Press.

Hearne, Robert R. and William K. Easter. 1997. The Economic and Financial Gains from Water Markets in Chile. *Agricultural Economics* 15: 187-199.

Hewitt, Julie A. and W. Michael Hanemann. 1995. A Discrete/Continuous Choice Approach to Residential Water Demand under Block Rate Pricing. *Land Economics* 71: 173-192.

Howe, Charles W., Dennis R. Schurmeier, and W. Douglass Shaw, Jr. 1986. Innovative Approaches to Water Allocation: The Potential for Water Markets. *Water Resources Research* 22: 439-445.

Kerr, Suzi and David Maré. 1997. Efficient regulation through tradeable permit markets: the United States lead phasedown. Working Paper 96-06, Department of Agricultural and Resource Economics, University of Maryland, College Park.

Kleit, Andrew N. and Dek Terrell. 2001. Measuring Potential Efficiency Gains from Deregulation of Electricity Generation: A Bayesian Approach. *Review of Economics and Statistics* 83: 523-30.

Krause, Kate, Janie M. Chermak, and David Brookshire. 2003. The Demand for Water: Consumer Response to Scarcity. *Journal of Regulatory Economics* 23: 167-191.

Lipsey, R.G. and Kelvin Lancaster. 1956. The General Theory of the Second Best. *Review of Economic Studies* 24: 11-32.

Mayer, Peter W., W. B. DeOreo, E. M. Opitz, J. C. Kiefer, W. Y. Davis, B. Dziegielewski, and J. O. Nelson. 1998. *Residential End-Uses of Water*. Denver: American Water Works Association Research Foundation.

Moffitt, Robert. 1990. The Econometrics of Kinked Budget Constraints. *Journal of Economic Perspectives* 4: 119-139.

Moffitt, Robert. 1986. The Econometrics of Piecewise-Linear Budget Constraints: A Survey and Exposition of the Maximum Likelihood Method. *Journal of Business and Economic Statistics* 4: 317-328.

Montgomery, W. David. 1972. Markets in Licenses and Efficient Pollution Control Programs. *Journal of Economic Theory* 5: 395-418.

National Climatic Data Center. 2007. Climate of 2007: June, U.S. Drought Watch. http://www.ncdc.noaaa.gov/oa/climate/research/2007/jun/us-drought.html. (accessed August 6, 2007).

Newell, Richard G. and Robert N. Stavins. 2003. Cost Heterogeneity and the Potential Savings from Market-Based Policies. *Journal of Regulatory Economics* 23: 43-59.

Newey, Whitney. 1987. Efficient Estimation of Limited Dependent Variable Models with Endogenous Explanatory Variables. *Journal of Econometrics* 36: 230-251.

Olmstead, Sheila M. Forthcoming. “Reduced-form vs. Structural Models of Water Demand Under Nonlinear Prices,” *Journal of Business and Economic Statistics*.

Olmstead, Sheila M., W. Michael Hanemann, and Robert N. Stavins. “Water Demand Under Alternative Price Structures,” *Journal of Environmental Economics and Management* 54: 181-198.

Parry, Ian W. H. and Antonio Bento. 2002. Estimating the Welfare Effect of Congestion Taxes: The Critical Importance of Other Distortions Within the Transport System. *Journal of Urban Economics* 51: 339-365.

Pels, Eric and Erik T. Verhof. 2004. The Economics of Airport Congestion Pricing. *Journal of Urban Economics* 55: 257-277.

Pigou, Arthur C. 1920. *The Economics of Welfare*. London: Macmillan and Co.

Pint, Ellen M. 1999. Household Responses to Increased Water Rates During the California Drought. *Land Economics* 75: 246-266.

Renwick, Mary E. and Sandra O. Archibald. 1998. Demand Side Management Policies for Residential Water Use: Who Bears the Conservation Burden? *Land Economics* 74: 343-359.

Rietveld, Piet, Jan Rouwendal, and Bert Zwart. 2000. Block Rate Pricing of Water in Indonesia: An Analysis of Welfare Effects. *Bulletin of Indonesian Economic Studies* 36: 73-92.

Sandmo, Agnar. 1975. Optimal Taxation in the Presence of Externalities. *Swedish Journal of Economics* 77: 87-98.

Small, Kenneth A. and Jia Yan. 2001. The Value of ‘Value Pricing’ of Roads: Second-Best Pricing and Product Differentiation. *Journal of Urban Economics* 49: 310-336.

Suen, Wing. 1990. Statistical Models of Consumer Behavior with Heterogeneous Values and Constraints. *Economic Inquiry* 28: 79-98.

Tietenberg, Tom. 1995. Tradable Permits for Pollution Control when Emission Location Matters: What Have We Learned? *Environmental and Resource Economics* 5: 95-113.

Timmins, Christopher. 2003. Demand-Side Technology Standards Under Inefficient Pricing Regimes: Are They Effective Water Conservation Tools in the Long Run? *Environmental and Resource Economics* 26: 107-124.

Weitzman, Martin L. 1973. Prices vs. Quantities. *Review of Economic Studies* 41: 477-491.

Weitzman, Martin L. 1977. Is the Price System or Rationing More Effective in Getting a Commodity to Those Who Need it Most? *The Bell Journal of Economics* 8: 517-524.

Table 1. Descriptive Statistics

Variable	Description	Units	Mean	Std. Dev.	Min.	Max.
w	Daily household water demand	kgal/day	.40	.58	0	9.78
	in season	kgal/day	.54	.71	0	9.78
	off season	kgal/day	.25	.34	0	7.16
w_{out}	Daily water demand outdoors	kgal/day	.22	.55	0	9.50
	in season	kgal/day	.36	.69	0	9.50
	off season	kgal/day	.07	.30	0	6.79
w_{in}	Daily water demand indoors	kgal/day	.17	.13	0	1.91
	in season	kgal/day	.17	.13	0	1.04
	off season	kgal/day	.17	.13	0	1.91
$P(w_{out}>0)$	Fraction obs. for which outdoor>0		.42	.49	0	1
$P(w_{in}>0)$	Fraction obs. for which indoor>0		>.99	.02	0	1
price	Observed marginal water price	$/kgal/mo	1.71	.57	.5	4.96
phat	Marginal water price instrument	$/kgal/mo	1.71	.53	.76	4.73
income	Gross annual household income	$000/yr	69.72	67.66	5.00	388.64
arid season	Irrigation season=1 / not=0		.51	.50	0	1
weath	Evapotransp. less effective rainfall	mm/day	5.06	8.42	-46.15	19.37
maxtemp	Maximum daily temperature	°C	24.12	8.78	0	42.78
famsize	Number of residents in household		2.78	1.34	1	9
bathrooms	Number of bathrooms in household		2.58	1.30	1	7
sqft	Area of home	000 ft^2	2.02	.82	.40	4.37
lotsize	Area of lot	000 ft^2	10.87	9.22	1.00	45.77
home age	Age of home	yrs/10	2.89	1.62	.07	5
evap cooling	Evaporative cooling=1 / not=0		.09	.28	0	1
region 1	Southern California		.37	.48	0	1
region 2	Arizona/Colorado		.28	.45	0	1
region 3	Northern		.26	.44	0	1
region 4	Florida		.09	.29	0	1

Notes: kgal is thousands of gallons.

Table 2. Model Testing Probability-Weighted Prices as Instruments

	DCC Estimates (OHS Paper)		Test Model 1		Test Model 2	
Variable	**Estimate**	**SE**	**Estimate**	**SE**	**Estimate**	**SE**
lnprice	-0.3408*	0.0298	-0.1751	0.1084	-0.3575*	0.0694
lnincome	0.1305*	0.0118	0.1421*	0.0319	0.1433*	0.0327
arid season	0.3070*	0.0247	0.3111*	0.0211	0.3156*	0.0209
weath	0.0079*	0.0013	0.0081*	0.0011	0.0078*	0.0010
maxtemp	0.0196*	0.0018	0.0193*	0.0015	0.0201*	0.0015
famsize	0.1961*	0.0056	0.1915*	0.0152	0.1949*	0.0156
bathrooms	0.0585*	0.0093	0.0481#	0.0250	0.0557*	0.0259
sqft	0.1257*	0.0140	0.1234*	0.0379	0.1349*	0.0393
lotsize	0.0065*	0.0009	0.0061*	0.0025	0.0077*	0.0025
home age	0.0867*	0.0219	0.0826	0.0592	0.0990	0.0613
home age^2	-0.0137*	0.0036	-0.0145	0.0098	-0.0171#	0.0102
evap cooling	0.2277*	0.0300	0.2329*	0.0823	0.2340*	0.0827
Fixed Effects	City-level		City-level		Region-level	
R^2 overall			0.20		0.19	
within			0.10		0.10	
between			0.36		0.32	

Notes: * significant at 5% (# at 10%). Dependent variable is natural log of daily household water demand (kgal). Model in column 1 is discrete-continuous choice model from Olmstead *et al.* (2007). Test models 1 and 2 are 2SLS random effects model for panel data, in which we instrument for marginal water prices. Estimates for city-level and region-level fixed effects are not reported. In all cases, N=25,668, with 1,082 households.

Table 3. Models of Indoor and Outdoor Water Demand

Variable	Indoor (annual)	Indoor (by season)	Outdoor (annual)	Outdoor (by season)
lnphat	-0.0727 (0.0577)	-0.0642 (0.0575)	-0.4457* (0.0452)	-0.4116* (0.0459)
lnphat wet season		-0.0349 (0.0286)		-0.1606* (0.0381)
perr			1.9777* (0.0720)	2.2680* (0.0963)
perr wet season				-0.6109* (0.1202)
lnincome	0.0669* (0.0279)	0.0671* (0.0278)	0.0939* (0.0196)	0.0950* (0.0196)
arid season	-0.0168 (0.0164)	-0.0341 (0.0215)	0.4267* (0.0211)	0.3388* (0.0282)
weath	0.0005 (0.0008)	0.0006 (0.0008)	0.0083* (0.0012)	0.0085* (0.0012)
maxtemp	-0.0012 (0.0012)	-0.0012 (0.0012)	0.0300* (0.0016)	0.0303* (0.0016)
famsize	0.2391* (0.0132)	0.2394* (0.0132)	0.0378* (0.0092)	0.0401* (0.0094)
bathrooms	0.0025 (0.0220)	0.0026 (0.0220)	0.0511* (0.0158)	0.0504* (0.0155)
sqft	0.0145 (0.0335)	0.0150 (0.0334)	0.1457* (0.0246)	0.1486* (0.0247)
lotsize	0.0031 (0.0022)	0.0031 (0.0021)	0.0104* (0.0015)	0.0104* (0.0015)
home age	0.0641 (0.0522)	0.0639 (0.0521)	0.0502 (0.0360)	0.0530 (0.0360)
home age^2	-0.0132 (0.0087)	-0.0131 (0.0086)	-0.0057 (0.0060)	-0.0061 (0.0060)
evap cooling	0.1456* (0.0704)	0.1454* (0.0703)	0.0829$^\#$ (0.0504)	0.0840$^\#$ (0.0501)

Notes: * significant at 5% ($^\#$ at 10%). For indoor models, dependent variable is natural log of daily household indoor demand and model is 2SLS random effects, with 25,136 observations. For outdoor models, dependent variable is daily household outdoor demand and model is 2SLS Tobit random effects, with 25,707 observations. About half of the observations fall in each season. Results for regional fixed effects not reported. The variable *perr* is the residual from the first stage (fitted price) equation.

Table 4. Summary of Elasticity Estimates

	Indoor Demand Elasticities	Outdoor Demand Elasticities
Overall	-0.0727 (0.0577)	-0.6836* (0.0693)
Arid Season	-0.0642 (0.0575)	-0.7365* (0.0893)
Wet Season	-0.0991 (0.0633)	-1.1750* (0.1073)

Notes: * significant at 5% (# at 10%). Elasticities are calculated for models reported in Table 3. Indoor elasticities are constant-elasticity demand model coefficients. Outdoor elasticities are estimated as follows, where α_{Tobit} is the Tobit price coefficient, and $\overline{P}$ and $\overline{w}$ are sample averages:

$$\varepsilon_{out} = \frac{\alpha_{Tobit} * \overline{P}(w_{out} > 0)}{\overline{w}_{out}}$$

Table 5. Price Elasticities of Demand, by Income/Lot Size Group

Household Group	Indoor Demand Elasticities	Outdoor Demand Elasticities
Rich, big lot	-0.1186 (0.0804)	-0.4837* (0.1028)
Poor, big lot	-0.0861 (0.1010)	-0.8017* (0.1233)
Rich, small lot	-0.0663 (0.0731)	-0.7758* (0.0946)
Poor, small lot	-0.0423 (0.0785)	-0.8690* (0.0938)

Notes: * significant at 5% ($^{\#}$ at 10%). The number of observations is 7,188 for rich, big lot; 4,016 for poor, big lot; 7,386 for rich, small lot; and 7,117 for poor, small lot.

Table 6. Shadow Prices, Market-clearing Prices Under Various Drought Policies

Drought Policy	Current Price Mean ($/kgal) [Std. Dev.]	Shadow Price Mean ($/kgal) [Std. Dev.]	Market-clearing Price Mean ($/kgal) [Std. Dev.]
Status quo (no drought policy)	1.71 [0.57]		
No outdoor watering		50.00 [0.00]	17.21 [14.50]
Outdoor watering once/week		6.92 [10.50]	6.86 [7.81]
Outdoor watering twice/week		4.98 [7.49]	3.90 [3.20]
Outdoor watering 3 times/week		3.55 [5.06]	2.77 [1.61]

Notes: All prices are for arid season only. We assume willingness-to-pay is at most $50 per thousand gallons.

Table 7. Estimated Welfare Impacts by Utility

	Shadow Price ($/kgal)			($/household/summer	
Utility	**Mean**	**Std. Dev.**	**Market Price ($/kgal)**	**DWL median and (90% C.I.)**	
Seattle, WA	2.91	0.41	2.92	2.3	[1.5, 3.9]
Eugene, OR	1.16	0.48	1.06	8.9	[5.5, 12.7]
Waterloo, ONT	2.25	0.68	2.10	7.3	[3.1, 16.1]
Cambridge, ONT	1.94	0.91	1.80	6.1	[2.1, 16.5]
Lompoc, CA	3.09	1.05	2.86	20.8	[12.7, 33.1]
Tampa, FL	2.04	1.19	1.96	16.5	[10.3, 28]
Highline, WA	3.74	1.79	3.30	21.5	[6.9, 57.4]
San Diego, CA	3.80	2.79	3.01	45.4	[26.7, 89.3]
Tempe, AZ	3.05	2.45	2.25	48.8	[26.9, 111.5]
Denver, CO	3.16	3.15	2.34	50.4	[29, 101.6]
Bellevue, WA	3.83	5.45	2.94	57.1	[12, 175.7]
Walnut Valley, CA	7.08	7.40	4.71	143.7	[74.5, 271.9]
Phoenix, AZ	5.02	5.87	3.85	110.5	[61.5, 222.1]
Scottsdale, AZ	7.43	7.10	5.07	152.1	[79, 364.8]
Northshore, WA	5.66	11.87	3.61	143.5	[8.3, 418.3]
Las Virgenes, CA	16.39	15.14	12.95	407.5	[221.6, 1018.3]
Sample-weighted Average	**4.98**	**4.35**	**3.90**	**92.2**	**[53.7, 196.6]**

Notes: Shadow prices and deadweight losses (DWL) are calculated for a two-day per week outdoor watering policy. Estimates and averages are for arid season only. DWL is estimated as the area under demand curves from 1500 nonparametric bootstrap replications. For indoor demand, $\widehat{w}_{in} = e^{Z\hat{\delta}} e^{H\hat{\beta}} p^{\hat{\alpha}} \tilde{y}^{\hat{\mu}}$. Let $C_{in} = e^{Z\hat{\delta}} e^{H\hat{\beta}} \tilde{y}^{\hat{\mu}}$. Then, $\widehat{w}_{in} = C_{in} p^{\hat{\alpha}}$ and the integral is: $C_{in} p(\frac{p^{\hat{\alpha}}}{(\widehat{\alpha}+1)}) = \frac{\widehat{w}_{in} p}{\widehat{\alpha}+1}$. For outdoor demand, $\widehat{w}_{out} = \widehat{\alpha} \ln p + \widehat{\mu} \ln \tilde{y} + \widehat{\delta} Z + \widehat{\beta} H$. Let $C_{out} = \widehat{\mu} \ln \tilde{y} + \widehat{\delta} Z + \widehat{\beta} H$. Then, $\widehat{w}_{out} = C_{out} + \widehat{\alpha} \ln p$, and the integral is: $C_{out} p - \widehat{\alpha} p + \widehat{\alpha} p \ln p = p(\widehat{w}_{out} - \widehat{\alpha})$.

Table 8. Average Deadweight Loss and Change in Surplus by Group

Group	Average DWL ($/arid season)		Average Change in Surplus		Average Change in Surplus/ Average Annual Income (1000s)	
Rich, big lot	$199	[112, 394]	-57	[-233, -11]	-0.5	[-1.9, -0.1]
Rich, small lot	64	[35, 171]	-48	[-75, -31]	-0.6	[-1, -0.4]
Poor, big lot	37	[22, 68]	-39	[-59, -27]	-1.4	[-2, -0.9]
Poor, small lot	41	[23, 84]	-26	[-40, -16]	-0.8	[-1.3, -0.5]
Average household	92	[54, 197]	-44	[-99, -27]	-0.6	[-1.4, -0.4]
Suppliers (per household)	-		136	[87, 290]		

Notes: Table reports medians and 90% confidence intervals from 1500 nonparametric bootstrap replications. DWL is welfare loss from the absence of a market. Average change in consumer and producer surplus results from introducing a constant, market-clearing price.

Table B.1. Robustness of Elasticity Estimates

	Indoor Demand Elasticities	Outdoor Demand Elasticities
City FEs	0.0881 (0.0883)	-0.3922* (0.1277)
Aggregate consumption	-0.0689 (0.0627)	-0.5959* (0.1029)

Notes: * significant at 5% (# at 10%). Results in Table B.1. should be considered relative to the estimates reported in Table 4.

Table B.2. Robustness of Heterogeneity Estimates

	City FE		Aggregate Data		City-specific Lot Size	
Household Sub-group	**Indoor**	**Outdoor**	**Indoor**	**Outdoor**	**Indoor**	**Outdoor**
Rich, big lot	0.0494 (0.1016)	-0.5183* (0.1573)	-0.0816 (0.0861)	-0.2943# (0.1483)	-0.1467 (0.0805)	-0.6883* (0.0988)
Poor, big lot	0.0819 (0.1250)	-0.6725* (0.1773)	-0.0512 (0.1023)	-0.8992* (0.1768)	-0.0234 (0.0915)	-0.7833* (0.1087)
Rich, small lot	0.0898 (0.0961)	-0.6727* (0.1559)	-0.0503 (0.0802)	-0.7085* (0.1391)	-0.0449 (0.0721)	-0.5979* (0.0930)
Poor, small lot	0.1296 (0.1079)	-0.6739* (0.1604)	-0.0836 (0.0802)	-0.7474* (0.1383)	-0.0809 (0.0825)	-0.8473* (0.0948)

Notes: * significant at 5% (# at 10%). Results in Table B.2. should be considered relative to the estimates reported in Table 5. In the city FE model for outdoor demand (column 2), the "rich, big lot" elasticity is weakly different from the others (p=.10).

Figure 1. Welfare loss from outdoor consumption restrictions

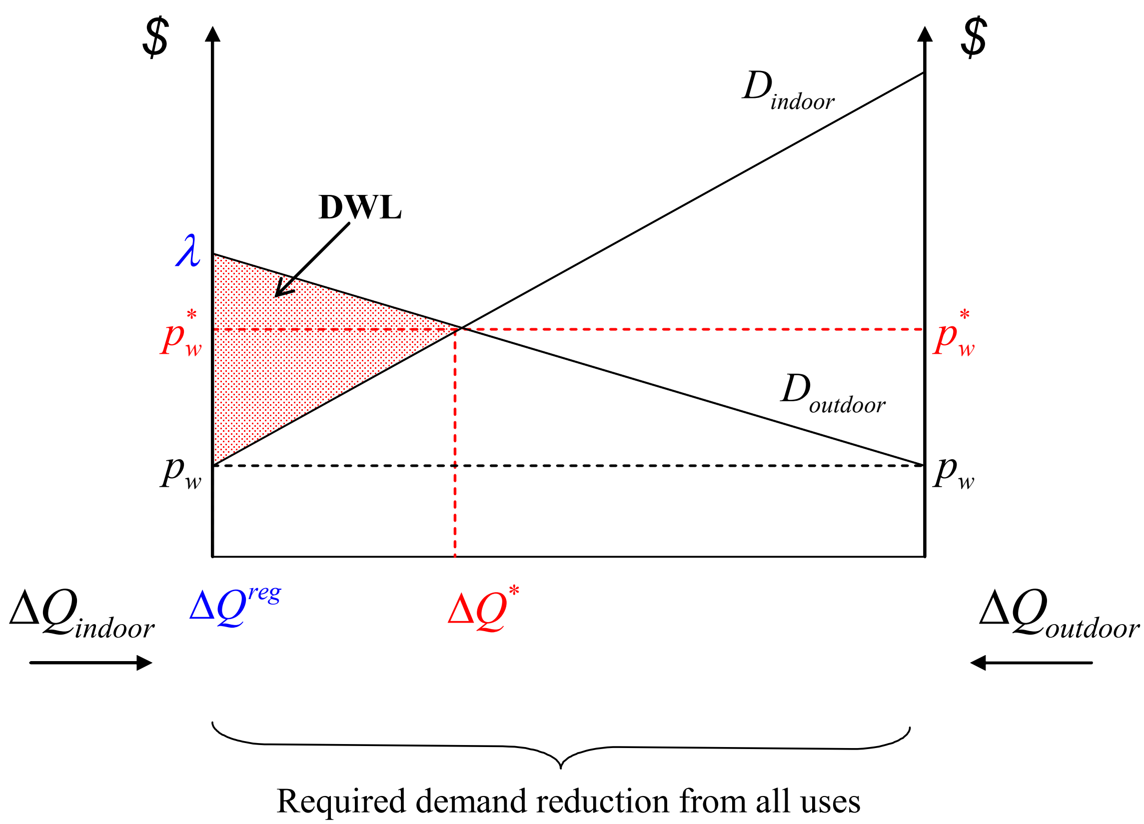

Figure 2. Welfare losses from outdoor consumption restrictions with heterogeneous outdoor demand

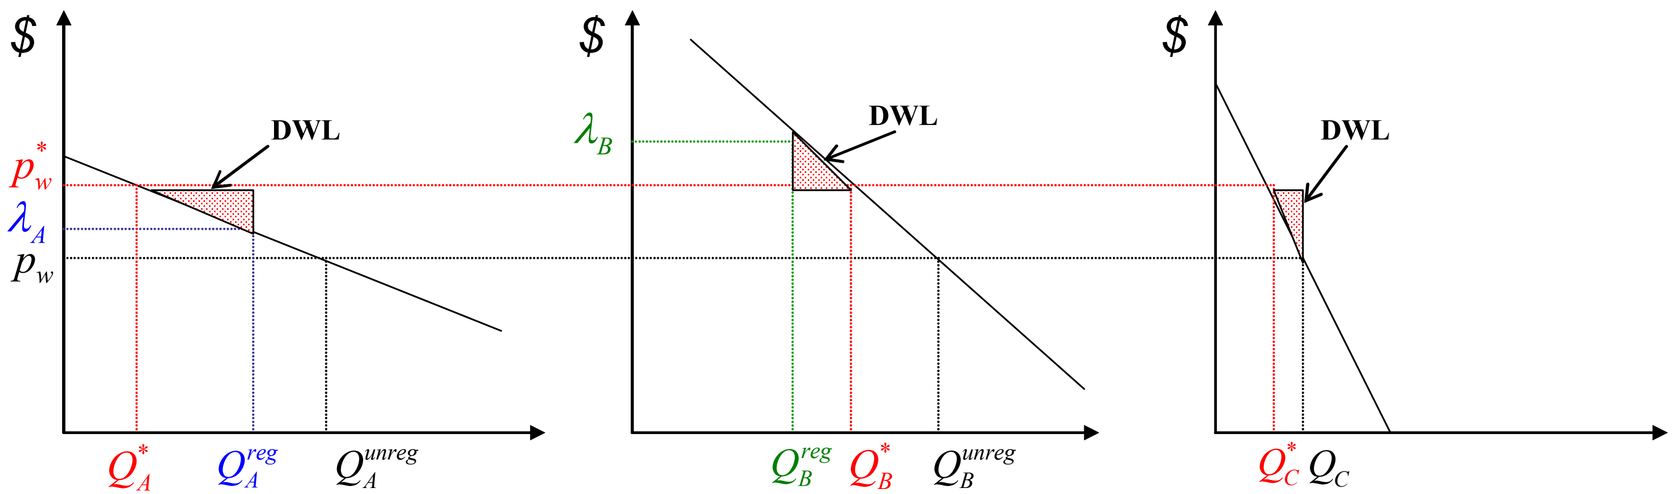

A
Relatively Elastic
Outdoor Demand

B
Somewhat Less Elastic
Outdoor Demand

C
Inelastic
Indoor Demand

(Where p_w^* is the market-clearing price for $Q_A^{reg} + Q_B^{reg} + Q_C = Q_A^* + Q_B^* + Q_C^*$).

Figure 3. Distribution of Shadow Prices

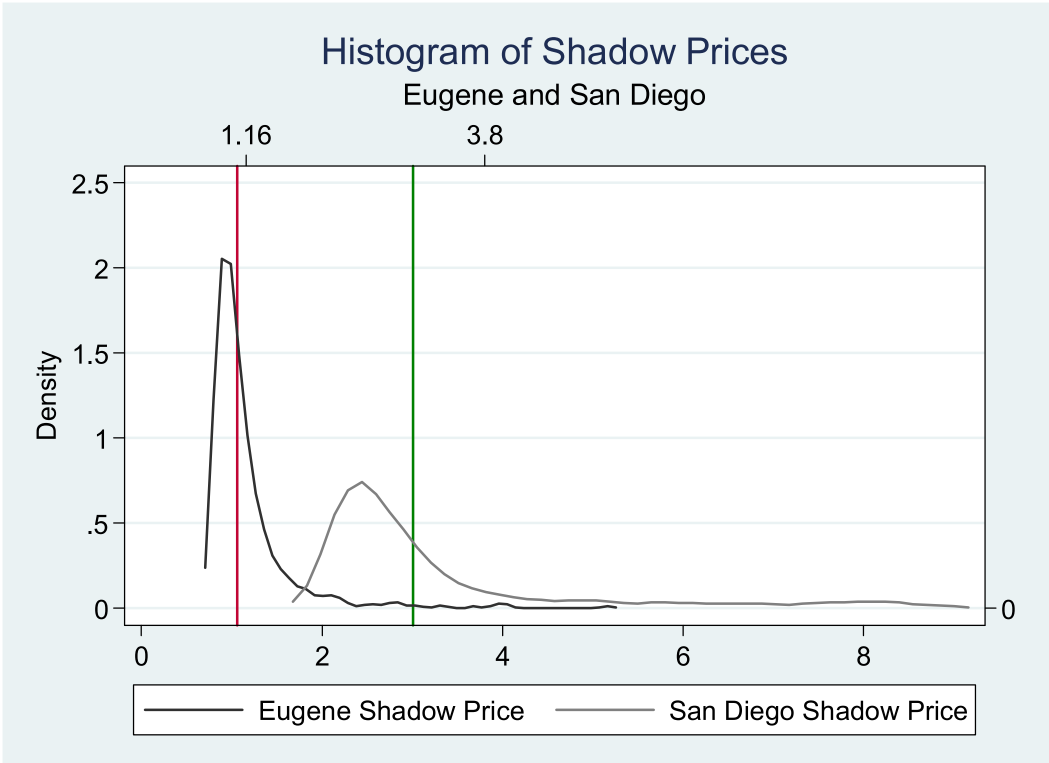

Notes: Distributions are of estimated shadow prices for Eugene, Oregon and San Diego, California, given a two days per week water policy. Numbers at the top of the figure identify average shadow prices for each city, and vertical lines represent the alternative policy of a market-clearing price, in dollars per thousand gallons.