

NBER WORKING PAPER SERIES

THE VALUE OF SCARCE WATER:
MEASURING THE INEFFICIENCY OF MUNICIPAL REGULATIONS

Erin T. Mansur
Sheila M. Olmstead

Working Paper 13513
<http://www.nber.org/papers/w13513>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
October 2007

For comments and guidance, we are grateful to Patrick Bayer, Nathaniel Keohane, Sharon Oster, Robert Stavins, Christopher Timmins, and seminar participants at Yale, Harvard, Dartmouth, Columbia, Camp Resources, the California Occasional Workshop on Environmental Economics, and the University of Maryland. The authors, alone, are responsible for any errors. The views expressed herein are those of the author(s) and do not necessarily reflect the views of the National Bureau of Economic Research.

© 2007 by Erin T. Mansur and Sheila M. Olmstead. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

The Value of Scarce Water: Measuring the Inefficiency of Municipal Regulations
Erin T. Mansur and Sheila M. Olmstead
NBER Working Paper No. 13513
October 2007
JEL No. L51,L95,Q25,Q28,Q58

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.

Erin T. Mansur
Yale School of Management
135 Prospect Street
P.O. Box 208200
New Haven, CT 06520-8200
and NBER
erin.mansur@yale.edu

Sheila M. Olmstead
Yale School of Forestry & Environmental Studies
Sage Hall, 205 Prospect Street
New Haven, CT 06511
sheila.olmstead@yale.edu

1. Introduction

Competitive markets are the most efficient avenues through which to allocate scarce resources to their highest-valued uses, under the standard assumptions. Yet, many important resources traditionally have not been managed through markets. Prominent examples include space on public roadways and airport runways, water supply, some telecommunications networks, electricity transmission and distribution, and, until recently, broadcast space on the electromagnetic spectrum and electricity generation. These goods and services have been heavily regulated, for a variety of reasons, with allocation and pricing (if any) controlled by public authorities. 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 switching from non-market to market-based regulation of municipal water supply during periods of drought.¹ During droughts, municipal water restrictions focus almost exclusively on the residential sector, rather than on commercial and industrial water users.² 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 the current approach. The welfare gains of a price-based approach could come from both substitution across uses within a household, and substitution across households.

¹ Between January 1997 and June 2007, moderate to extreme drought conditions affected, on average, 20 percent of the contiguous United States. This 10-year average masks a great deal of variation; in June 2007, 41 percent of the contiguous United States was experiencing moderate to extreme drought (National Climatic Data Center 2007).

² Of municipal water consumption, one-half to two-thirds is residential (Dixon *et al.* 1996, Gleick *et al.* 2003).

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. Our econometric approach presents a substantial challenge, in that: (1) sixty percent of sample households face increasing-block prices, creating piecewise linear budget constraints under which demand and marginal price are systematically correlated; and (2) outdoor water use is censored at zero – households do not use water outdoors every day.

We use these estimates to determine whether current policies target the water uses that households, themselves, would reduce in response to price increases. We find that outdoor watering restrictions do mimic household reactions to price increases, on average, as outdoor use 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 the marginal unit of water may vary greatly among households.

We divide households into four groups, based on income and lot size, and estimate separate end-use elasticities by group to assess the degree of household heterogeneity in these markets. We estimate shadow prices for water among all customers under four drought policy scenarios of increasing stringency. We then estimate utility-level market-clearing prices for these drought policy scenarios. Under a typical drought policy, one that limits outdoor watering to two days per week, estimated shadow prices are, on average, 178 percent higher than the current average marginal water price in these markets, while market-clearing prices are 118 percent higher than current prices.

Note that the analysis does not consider whether water prices, in general, are set at efficient levels – only whether the particular 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.

However, our results have important implications for municipal policies. The welfare gains from allowing substitution across uses within the average household during an average summer drought amount to about \$27. Allowing substitution across households, as well as within, welfare gains from a price-based approach are approximately \$81 per household, about 25 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 also would have consequences for allocation. Drought prices would enable those customers who are the least price sensitive, the wealthy consumers with large lots, to reduce consumption much less than poor 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 \$158 per customer. Households would be worse off by \$77, on average. Rebates to consumers could make everyone better off.

The paper proceeds as follows. In Section 2, we discuss some simple theory and the related literature. Section 3 examines the econometric models, and Section 4 presents the data. In Section 5, we show our price elasticity estimates for various 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). Studies have produced theoretical and empirical estimates of the gains from increasing the influence of markets on: the distribution of space in the broadcast spectrum (Melody 1980, McMillan 1994); traffic congestion on roadways (Winston 1985, Kraus 1989, 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 resource allocation in these cases derive largely from heterogeneity in marginal benefits across potential consumers.³

³ 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 market-based policies 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, and its close associate, ability to pay, may seem to be unjust allocation criteria.

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, lawns, hot tubs). 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 ensures that a citywide required demand reduction will be achieved by 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 own preferences and ability to pay (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 demand curves for indoor and outdoor water use against a required demand reduction (on the horizontal axis).⁴ 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 (\bar{P}). Under a market-clearing price (P^*), some of the citywide required reduction would take place indoors, and this wedge would disappear. All else equal, the less elastic indoor demand, the smaller the welfare losses will be from the current approach.

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. Putting this all together, the DWL from non-market regulation has three components: some households consume too much outdoors, some consume too little outdoors, and everyone consumes too much indoors.

⁴ We assume that supply, in the short-run situation of drought, is perfectly inelastic.

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.

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. One experimental economics 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.

A few studies have empirically estimated the impacts of non-price conservation programs on aggregate residential water demand, sometimes jointly with price elasticities (Michelsen *et al.* 1998, Corral 1997). Such studies analyze city-level demand across cities, constructing indices of the relative stringency of residential demand management programs. This approach impedes the

direct comparison of the costs of specific non-price programs with hypothetical price increases to achieve the required level of reduction in aggregate water use.⁵ We know of only one case in which such a comparison has been made. Timmins (2003) compares a mandatory low-flow appliance regulation with a modest water tax, using aggregate consumption data from 13 groundwater-dependent California cities. Under all but the least realistic of assumptions, he finds the tax to be more cost-effective than the technology standard in reducing groundwater aquifer lift-height in the long run.

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.⁶ One end-use water demand study has been performed, including estimates of price elasticity for specific uses of water (Mayer *et al.* 1998). This study was the first to use data that reliably separate household water consumption into its component end uses. We use the same data, but ask different questions, and account for price endogeneity.⁷

Most U.S. water utilities charge non-linear tariffs to residential customers. These price schedules typically require a structural approach to obtain unbiased estimates of price elasticity

⁵ Not all studies of non-price demand management are able to detect an impact of such programs on residential demand. One study of draconian outdoor watering restrictions in Corpus Christi, Texas, during a drought of record in summer 1996 found no impact of the restrictions on residential consumption (Schultz *et al.* 1997).

⁶ Meta-analyses suggest 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 (Espey *et al.* 1997, Dalhuisen *et al.* 2004).

⁷In estimating demand elasticities, Mayer *et al.* (1998) use average prices, not marginal prices. Ours is the first application of end-use demand estimation to treat non-linear prices in a simultaneous equations framework.

and other model parameters. In the case of increasing-block prices, marginal price and the quantity consumed are positively correlated. 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 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). We use the parameter estimates from Olmstead *et al.* (2007) to construct our price instruments, building on their approach to ask and answer policy-relevant questions regarding droughts, non-price demand management, and hypothetical price-based approaches.

3. Econometric Models

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.⁸ 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, since we know those estimates to be unbiased.⁹

The equation for this total demand function is (1), 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.¹⁰ The error structure comprises θ , a household heterogeneity parameter, and ν , the residual.

$$\ln w_{total_{it}} = \alpha \ln p_{it} + \mu \ln \tilde{Y}_i + \delta Z_{it} + \beta H_i + \theta_i + \nu_{it} . \quad (1)$$

Or, in matrix notation:

$$\ln \mathbf{w}_{total} = \mathbf{X}' \boldsymbol{\gamma}_{total} + \mathbf{v}_{total} \quad (2)$$

⁸ 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.

⁹ 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.

¹⁰ 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).

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 (2), except for the different dependent variable, is represented in (3).

$$\ln \mathbf{w}_{in} = \mathbf{X}'\boldsymbol{\gamma}_{in} + \mathbf{v}_{in} \quad (3)$$

For outdoor demand, we estimate a Tobit model, described in (4).

$$\begin{aligned} \mathbf{w}_{out}^* &= \mathbf{X}'\boldsymbol{\gamma}_{out} + \mathbf{v}_{out} \\ w_{out_{it}} &= \begin{cases} w_{out_{it}}^* & \text{if } w_{out_{it}}^* > 0 \\ 0 & \text{otherwise} \end{cases} \end{aligned} \quad (4)$$

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

Daily demand data are drawn from 1,082 households, randomly selected from billing databases in 11 urban areas across the United States and Canada, served by 16 water utilities. Observed households are detached, single-family homes, with no apartments or other multi-family housing in the sample.¹¹ Table 1 provides descriptive statistics.

Households were observed for two periods of two weeks each, once in an arid season, and once in a wet season. Total demand was disaggregated into its indoor and outdoor

¹¹ Households were randomly sampled from single-family homes within utility customer databases. Sampling procedures, response rates, and statistical tests for selection are described in Mayer *et al.* (1998), Appendix A.

components 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.¹² Leaks and unknown uses are included in total demand, but are not modeled explicitly as either indoor or outdoor consumption.¹³

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.

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 ft², 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

¹² 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.

¹³ Leaks comprise approximately 6 percent of total consumption, and unknown uses approximately 1 percent. 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.

consumption as follows: rich, big lot (43 percent); rich, small lot (23 percent); poor, big lot (15 percent); poor, small lot (19 percent).

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.¹⁴ Given cross-sectional variation and changes over time, the data contain 26 price structures and 47 different marginal prices, ranging from zero to just under \$5 per thousand gallons. 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.

Price variation in the sample is primarily in the city cross-section. If we regress the natural log of observed marginal prices on our set of regional dummies, we obtain an R-squared of 0.34. 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. Table 1 demonstrates that the mean of our price instrument, \hat{p} , is equal to the mean of observed marginal water prices, and the standard deviation is slightly smaller.¹⁵ Data on household characteristics, including information on annual income, family size, square footage and age of homes, lot size, number of bathrooms, and the presence of evaporative cooling were gathered by survey.¹⁶ We control for season (arid vs. wet), as well as daily weather variables, including maximum daily temperature, and

¹⁴ 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).

¹⁵ 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 should 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.

¹⁶ 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.

evapotranspiration less effective rainfall (0.6 times total rainfall). Finally, we construct regional fixed effects to control for long-run climate variation not absorbed by the daily and seasonal weather variables.¹⁷

5. Results

5.1. End-Use Price Elasticity

We begin by estimating a total water demand model as in (2), 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, unsurprisingly, somewhat 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.

¹⁷ Regions are as follows: (1) Southern California (Las Virgenes Municipal Water District, City of San Diego, Walnut Valley Water District, and City of Lompoc); (2) Arizona and Colorado (Phoenix, Tempe, Scottsdale, and 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); and (4) Florida (City of Tampa Utilities).

We then separate total demand into indoor and outdoor consumption and estimate the models given in (3) and (4). In Table 3, we report 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.¹⁸ 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). In fact, essentially all of the strong seasonal variation we observe in total water consumption is attributable to outdoor use.¹⁹

Results reported in Table 4 might suggest that regulators' focus on outdoor consumption as a target of command-and-control water conservation policies provides a good first approximation to a price-based approach. Indeed, outdoor uses are the uses that households,

¹⁸ 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.

¹⁹ 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.

themselves, would choose to cut back the most in response to a price increase. But an important cost of the prescriptive approach derives from the heterogeneity of regulated entities – in this case, households. If households are heterogeneous in their preferences for outdoor consumption, across-the-board outdoor water use restrictions will ignore that heterogeneity, generating welfare losses relative to a price-based approach.

5.2. Consumer 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.²⁰ 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 (-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.²¹

5.3. Robustness

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

²⁰ 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).

²¹ 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.

aggregate seasonal household demand on the covariates of interest. We find our results to be qualitatively robust to these alternative models and others. See Appendix B for details.

6. Simulations and Discussion

In this section, we describe our drought policy simulations and estimate shadow and market-clearing prices. We then examine the potential welfare gains from and distributional implications of using prices to regulate municipal drought.

6.1. Simulating Drought Policy

Traditional regulations typically limit the number of days in a week that households may use water outdoors, whether for watering lawns, washing automobiles, or filling swimming pools. During the most recent multi-year, statewide drought in California, 1987-1992, between 65 and 80 percent of urban water utilities implemented outdoor watering restrictions. Most of these were mandatory, but not well enforced (Dixon *et al.* 1996). The stringency of these type-of-use restrictions (and their enforcement) varies greatly. 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 these shadow prices (λ_{it}), 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 (4) – to obtain their willingness-to-pay for the marginal unit of water.²² 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

²² 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). So 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 (5). 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:

$$E(\Delta w_{out_{it}}) = \begin{cases} (prob_{it} - \frac{2}{7}) * w_{out_{it}}^{cond} & \text{if } prob_{it} > \frac{2}{7} \\ 0 & \text{otherwise} \end{cases} \quad (5)$$

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

Caveats

We face a choice of simulating utilities' stated policies, or simulating the actual results of their stated policies, and we choose the former. This gives rise to two important issues. First, our analysis assumes that the necessary aggregate (utility-wide) demand reduction is that which would be achieved under full compliance with the various drought policies. We could easily add a probability of compliance that is less than one to the analysis, but doing so would be equivalent to simulating different drought policies than those we select.²³ It is likely that compliance will be higher under a price-based policy than under the current prescriptive approach.²⁴ Nevertheless,

²³ If non-compliance is punished with a fine, some of the benefits of drought pricing may be realized under the command-and-control policy, depending on the magnitude of the fine.

²⁴ "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.

were utilities to estimate market-clearing prices, they might want to charge slightly higher prices, anticipating less than full compliance, depending on the importance of meeting the demand reduction.²⁵

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

²⁵ 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 transactions costs involved in establishing a household-level trading regime would likely be prohibitive.

increase consumption by less than one percent.²⁶ 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.

6.2. 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. For example, a price-based approach would lead to smaller reductions in outdoor consumption by the least elastic groups, larger reductions by the most elastic groups, and, perhaps, small reductions in indoor use by all groups (Figure 2). 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 price consumers actually pay (\$1.79).

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

²⁶ 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.

establishing a price-based policy. For example, Figure 3 provides histograms of shadow prices in two cities given a policy limiting outdoor watering to two days a week.²⁷ Even in cities with relatively small standard deviations in shadow prices (like Eugene, Oregon), we see there is a lot of variation and that 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; all gains from trade will be realized.

6.3. 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 $E(\Delta w_{out,i})$ in (5), for all households within a utility. From the households' perspective, these savings could be achieved indoors, outdoors, or by purchasing "credits" from another household.^{28, 29} We then identify the market-clearing price for that aggregate reduction, constraining households to non-negative consumption.³⁰

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

²⁷ We chose two cities (Eugene, Oregon, and San Diego, California) that provide examples with low and average amounts of variation in shadow prices, respectively.

²⁸ 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.

²⁹ 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.

³⁰ Market-clearing prices are estimated for the irrigation season, assuming that drought regulations are implemented primarily during arid months.

and magnified by their variation. In Eugene, where there was relatively little shadow price 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 across all cities. Like the shadow prices, market-clearing prices increase monotonically with the stringency of watering restrictions. Within a utility, there is a common price. While not as large as the variation within shadow prices, we still find that market-clearing prices vary substantially 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.

6.4. 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.³¹

For each utility, we simulate the welfare effects of a two-day per week watering policy over a 180-day drought-struck irrigation season. Deadweight losses (DWL) under the current regime represent the estimated benefit to the average household of introducing drought pricing.³² First, if we assume households are homogeneous and allow them to substitute some indoor for outdoor reductions, the average DWL is \$26.90 per household.³³ Next, we allow substitution across households. Table 7 summarizes our findings on a per household basis. Estimated DWL ranges from \$1.71 per household, in the service area of Seattle Public Utilities, to \$437 per

³¹ We estimate deadweight loss by integrating demand curves, as discussed in Table 7.

³² 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.

³³ This average DWL is estimated using the coefficients reported in Table 3.

household, in the service area of 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 about \$81 per household through the introduction of drought pricing.

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 very large magnitudes. 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 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). Applying a price-based approach to drought policy 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 price-based drought management results in more households paying something closer to marginal social cost, the pre-existing distortion does not

change the basic nature of our results.³⁴ 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 area for further research, but it is beyond the scope of this analysis.

6.5. 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 of shifting regulatory approaches depend on the allocation of property rights. Prices would re-distribute scarce water such 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.

Figure 4 describes the movement of water consumption across income/lot size groups that would be brought about by the shift from prescriptive to price-based allocation of water during periods of scarcity. For the purpose of exposition, we choose a two-day per week watering policy, and the market-clearing price that would generate the equivalent level of aggregate demand reduction. The consumption share of rich, big lot households would grow substantially, primarily as a result of consumption reduction by poor, small lot households, and to a smaller extent by the other groups. 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,

³⁴ 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).

would fall from 23 to 16 percent, with smaller reductions in consumption shares by the remaining two groups. Absolute consumption falls quite drastically among all groups under both types of drought policies.

Given the changes in consumption shares induced by moving from the current approach to drought pricing, 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, \$158 per household, would be used for this purpose, as utilities in the United States are usually restricted to zero (or very small) profits.³⁵ Households' minimum willingness to accept payment to support the policy would equal their average change in CS. Within lot size, a price-based system without rebates would be regressive.

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 recaptured DWL. Market-based policies need not be regressive.

³⁵ 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 traditional municipal drought policies. 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 water under conditions of scarcity.

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. While we find that there is some potential for substitution within households across end-uses, most of the welfare gains from a price-based approach result from the variation across households in values for scarce water.

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.³⁶ 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.

³⁶ 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, drought or no drought, 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, both 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.³⁷

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

Let $\underline{w}_k^*(.) = e^{Z\delta} e^{X\beta} p_k^\alpha \tilde{Y}_k^\mu$, or optimal consumption in block k . Then, conditional 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 conditional price as in (A.3).

³⁷ 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).

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

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

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 \hat{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.

$$\hat{p} = \Pr A * p_1 + \Pr B * (.5p_1 + .5p_2) + \Pr C * p_2 \quad (\text{A.4})$$

Where:

$$\Pr A = \int_0^{\frac{w_1}{\underline{w}_1^*(.)}} f(e^\eta) de^\eta \quad (\text{A.5})$$

$$\Pr C = \int_{\frac{w_1}{\underline{w}_2^*(.)}}^{\infty} f(e^\eta) de^\eta$$

and $\Pr B = 1 - \Pr A - \Pr C$

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 (recall that about 94 percent of variation in sample prices can be explained by household FEs, alone).

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, as we have done. 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 the 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 their daily counterparts. 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.
- 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.
- Corral, Leonardo R. 1997. Price and Non-Price Influence in Urban Water Conservation. Ph.D. diss. University of California at Berkeley.
- 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.
- Espey, Molly, James A. Espey, and W. Douglass Shaw. 1997. Price Elasticity of Residential Demand for Water: A Meta-analysis. *Water Resources Research* 33: 1369-1374.
- 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.
- Gleick, Peter H., Dana Haasz, Christine Henges-Jeck, Veena Srinivasan, Gary Wolff, Katherine Kao Cushing, and Amardip Mann. 2003. *Waste Not, Want Not: The Potential for Urban Water Conservation in California*. Oakland, CA: Pacific Institute.
- 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.

- 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.
- Kraus, Marvin. 1989. The Welfare Gains from Pricing Road Congestion Using Automatic Vehicle Identification and On-Vehicle Meters. *Journal of Urban Economics* 25: 261-281.
- 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.
- McMillan, John. 1994. Selling Spectrum Rights. *Journal of Economic Perspectives* 8: 145-162.
- Melody, William H. 1980. Radio Spectrum Allocation: The Role of the Market. *American Economic Review, Papers and Proceedings* 70: 393-397.
- Michelsen, Ari M., J. Thomas McGuckin, and Donna M. Stumpf. 1998. *Effectiveness of Residential Water Conservation Price and Nonprice Programs*. 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.noaa.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 Non-linear Prices,” *Journal of Business and Economic Statistics*.
- Olmstead, Sheila M., W. Michael Hanemann, and Robert N. Stavins. Forthcoming. “Water Demand Under Alternative Price Structures,” *Journal of Environmental Economics and Management*.
- 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.
- Schultz, Martin T., Sheila M. Cavanagh, Bisheng Gu, and David J. Eaton. 1997. The Consequences of Water Consumption Restrictions During the Corpus Christi Drought of 1996. Draft Report, LBJ School of Public Affairs, University of Texas, Austin.
- 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.
- Winston, Clifford. 1985. Conceptual Developments in the Economics of Transportation: An Interpretive Survey. *Journal of Economic Literature* 23: 57-94.

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 \bar{P} and \bar{W} are sample averages:

$$\epsilon_{out} = \frac{\alpha_{Tobit} * \bar{P}(w_{out} > 0)}{\bar{W}_{out}}$$

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)	1.79 [0.65]		
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

Utility	Shadow Price (\$/kgal)		Market Price (\$/kgal)	DWL (\$/household/summer)
	Mean	Std. Dev.		
Seattle, WA	2.91	0.41	2.92	1.71
Eugene, OR	1.16	0.48	1.06	2.93
Waterloo, ONT	2.25	0.68	2.10	3.48
Cambridge, ONT	1.94	0.91	1.80	4.31
Lompoc, CA	3.09	1.05	2.86	5.61
Tampa, FL	2.04	1.19	1.96	10.08
Highline, WA	3.74	1.79	3.30	12.88
San Diego, CA	3.80	2.79	3.01	26.04
Tempe, AZ	3.05	2.45	2.25	26.96
Denver, CO	3.16	3.15	2.34	29.93
Bellevue, WA	3.83	5.45	2.94	58.33
Walnut Valley, CA	7.08	7.40	4.71	94.85
Phoenix, AZ	5.02	5.87	3.85	106.54
Scottsdale, AZ	7.43	7.10	5.07	113.04
Northshore, WA	5.66	11.87	3.61	234.56
Las Virgenes, CA	16.39	15.14	12.95	437.48
Sample-weighted Average	4.98	4.35	3.90	80.63

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. For indoor demand,

$\hat{w}_{in} = e^{z\hat{\delta}} e^{H\hat{\beta}} p^{\hat{\alpha}} \tilde{Y}^{\hat{\mu}}$. Let $C_{in} = e^{z\hat{\delta}} e^{H\hat{\beta}} \tilde{Y}^{\hat{\mu}}$. Then, $\hat{w}_{in} = C_{in} p^{\hat{\alpha}}$ and the integral is: $C_{in} p \left(\frac{p^{\hat{\alpha}}}{(\hat{\alpha} + 1)} \right) = \frac{\hat{w}_{in} p}{\hat{\alpha} + 1}$. For outdoor

demand, $\hat{w}_{out} = \hat{\alpha} \ln p + \hat{\mu} \ln \tilde{Y} + \hat{\delta} Z + \hat{\beta} H$. Let $C_{out} = \hat{\mu} \ln \tilde{Y} + \hat{\delta} Z + \hat{\beta} H$. Then, $\hat{w}_{out} = C_{out} + \hat{\alpha} \ln p$, and the integral is:

$C_{out} p - \hat{\alpha} p + \hat{\alpha} p \ln p$.

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
Rich, big lot	\$223.9	-\$73.2	-.0006
Rich, small lot	24.2	-109.3	-.0014
Poor, big lot	11.7	-75.6	-.0026
Poor, small lot	30.6	-47.2	-.0015
Average household	80.6	-77.1	-.0011
Suppliers (per household)	-	157.7	

Notes: 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

Household Sub-group	City FE		Aggregate Data		City-specific Lot Size	
	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


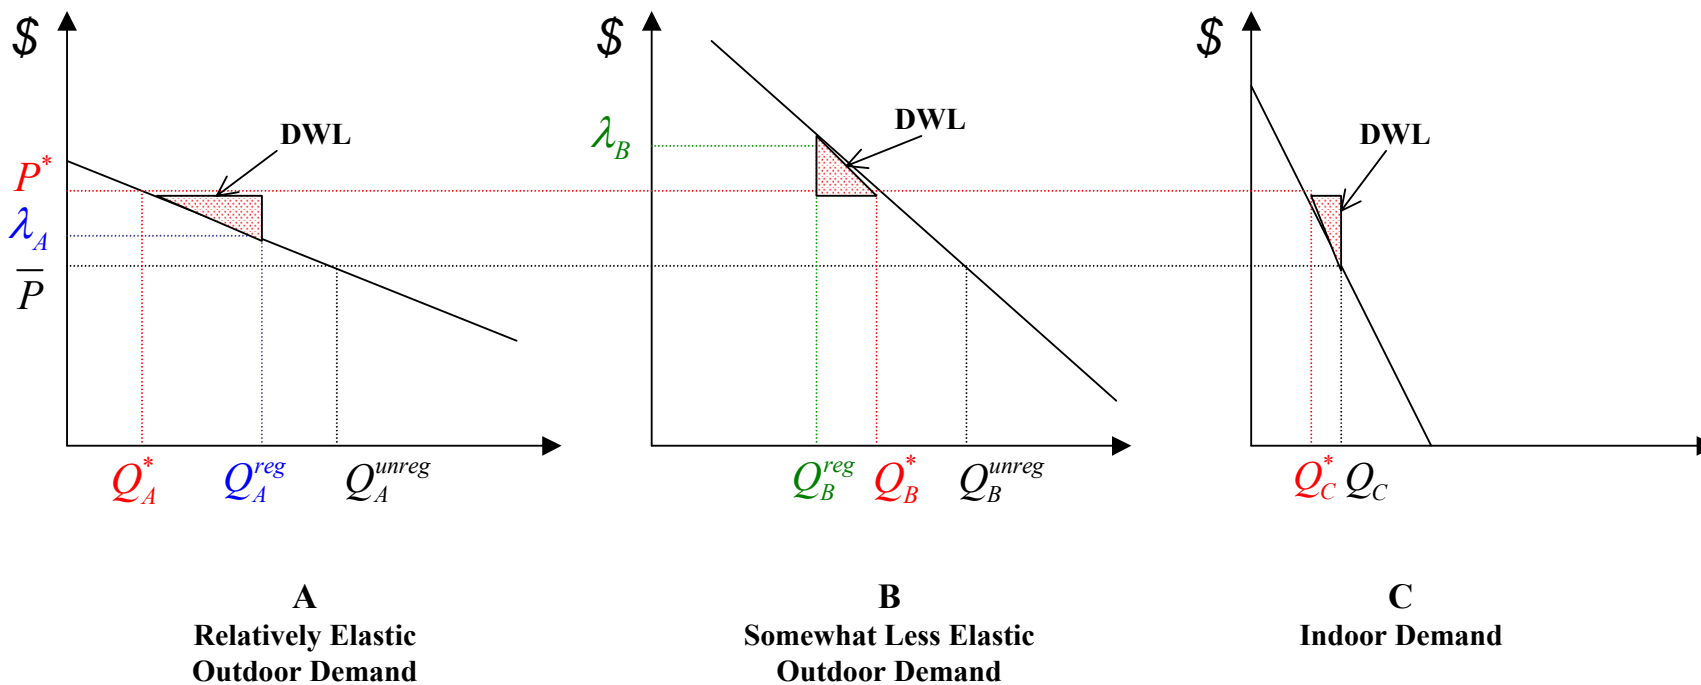



Figure 2. Welfare Losses from Outdoor Consumption Restrictions with Heterogeneous Outdoor Demand



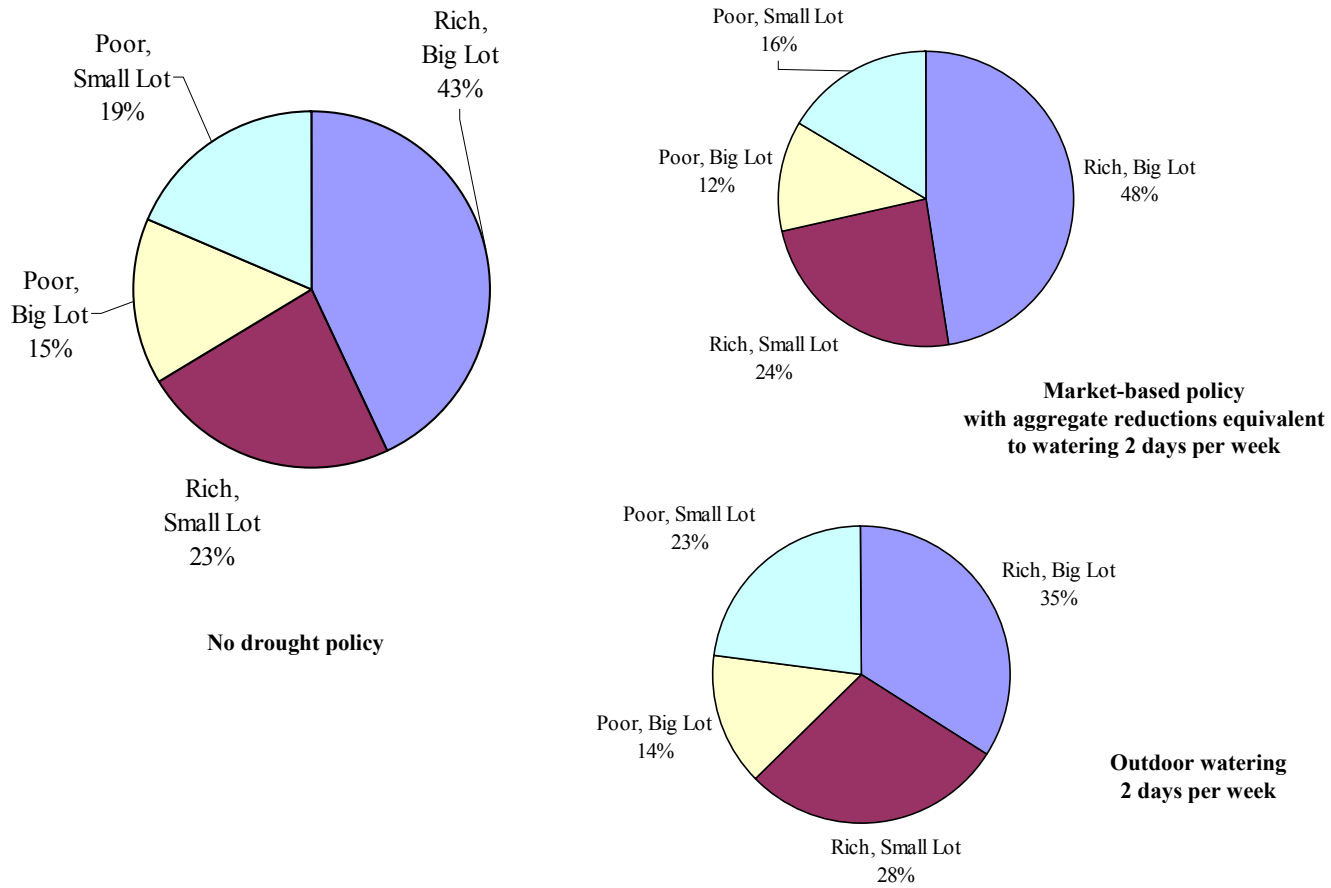
(Where P^* is the market-clearing price for $Q_A^{reg} + Q_B^{reg} + Q_C = Q_A^* + Q_B^* + Q_C^*$).

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.

Figure 4. Allocation of Water Consumption by Drought Policy



Notes: Total arid season consumption (the “size of the pie”) is 6,815 kgal with no drought policy, and 4,607 kgal with either drought policy.