Spatial Externalities in Groundwater Extraction: Evidence from California Agriculture

Fiona Burlig, Louis Preonas, and Matt Woerman*

December 24, 2018

Preliminary and Incomplete
Not for citation or circulation

Abstract

Groundwater is a common-pool resource essential for agricultural production. When farmers extract a marginal unit of groundwater, this lowers nearby groundwater levels and increases their neighbors' groundwater pumping costs. This paper estimates farmers' elasticity of demand for groundwater, in order to empirically investigate the magnitude of this spatial "pumping cost" externality. We assemble a novel dataset that combines (i) detailed microdata on farmers' electricity consumption, (ii) rich data from technical audits of these farmers' pump efficiencies, and (iii) publicly available measurements of groundwater depths in California aquifers. Using exogenous variation in electricity prices, we estimate farmers' price elasticities of demand for both electricity (-1.17) and groundwater (-1.12) to be much larger than previous estimates in the literature. We then calculate the extent to which each farm lowers its neighbors' economic surplus by removing water from their shared aquifer. Our preliminary results suggest that the magnitude of the "pumping cost" externality is likely smaller than farmers' private costs of groundwater pumping.

Keywords: Common-pool resources, groundwater, agriculture, electricity
JEL Codes: Q15, Q25, Q41

*Burlig: Harris School of Public Policy and Energy Policy Institute, University of Chicago. 1155 E 60th St, Chicago, IL 60637. Email: burlig@uchicago.edu. Preonas: Department of Economics and Energy Policy Institute, University of Chicago. 5757 S University Avenue, Chicago, IL 60637. Email: lpreonas@uchicago.edu. Woerman: Department of Agricultural and Resource Economics, University of California Berkeley. 207 Giannini Hall, Berkeley, CA 94720. Email: matt.woerman@berkeley.edu. We thank Amy Ando, Maximilian Auffhammer, Kathy Baylis, Michael Greenstone, Koichiro Ito, Katrina Jessoe, Dave McLaughlin, Andrew Stevens, Catherine Wolfram, and seminar participants at UC Energy Camp, Mississippi State, UC Davis, and University of Illinois for helpful comments and suggestions. Chen Sui, Yixin Sun, and Xinyi Wang provided excellent research assistance. Karen Notsund provided invaluable support in obtaining data. We received generous funding in support of this project from the Giannini Foundation of Agricultural Economics and the Sloan Foundation (The E2e Project). All remaining errors are our own.

1 Introduction

Groundwater is a classic example of a common-pool resource (Ostrom (1990)). When agent i makes a private decision to extract water from an underground aquifer, this lowers the depth of the water below the surface. If many agents share the same aquifer, then agent i's decision to extract imposes an externality on all other agents—who must now expend more energy to extract groundwater from greater depths. In theory, a social planner could increase welfare by taxing agent i's marginal unit of groundwater extraction, where the size of the tax would be calibrated to the marginal decrease in surplus summed across all other agents.

In this paper, we investigate the size of this pumping cost externality for California farmers. California produces 17 percent of U.S. crop value, and farmers rely heavily on groundwater for irrigation. With thousands of farmers extracting water from a few large aquifers, an individual farmer's extraction likely increases pumping costs for thousands of other farmers. Moreover, groundwater pumping is largely unregulated in California, and most farmers face no meaningful restrictions on the intensive margin of pumping. Hence, these farmers have little incentive to internalize the open-access externality, whereby their own marginal pumping decisions reduce their neighbors' economic surplus from groundwater consumption.

Estimating the size of this externality requires estimating the price elasticity of demand for agricultural groundwater. This has proven difficult, in part because groundwater use is typically neither priced nor measured. However, we leverage the fact that electricity is the main variable input in groundwater extraction. Given data on electricity prices and quantities, the relationship between energy use and water use for each farm, and groundwater levels, we are able to construct accurate measures of groundwater prices and quantities. We assemble to a novel dataset that combines (i) confidential electricity consumption data for the universe of agricultural customers served by Pacific Gas & Electric (PGE), California's largest electric utility; (ii) technical pump efficiency audits for nearly 12,000 individual groundwater extraction points; and (iii) publicly available measurements of groundwater depths over time for all major California aquifers. By leveraging exogenous variation in both electricity prices—coming from changes in PGE's tariffs schedules—and average groundwater levels, we causally estimate farmers' elasticities of demand for both electricity and groundwater.

We first estimate the price elasticity of demand for electricity in the agricultural sector. We find an elasticity estimate of -1.17, which is much larger than prior estimates of the price elasticity of electricity demand in both the residential and commercial/industrial sectors. Next, we estimate the price elasticity of demand for groundwater, where we separately identify the effect of changes to farmers' (effective) price of groundwater coming from variation in electricity prices vs. variation in groundwater depths. We recover nearly identical groundwater demand elasticity estimates: -1.39 for electricity-induced price changes, and -1.37 for depth-induced price changes. These statistically indistinguishable estimates suggest that farmers are equally attentive and responsive to either source of variation in groundwater pumping costs, consistent with the predictions of standard Neoclassical theory. We also estimate a single elasticity for demand for groundwater of -1.12, identified using *only* changes in PGE's agricultural electricity tariffs. These estimates are again much larger than most groundwater demand elasticities from the existing literature, which may reflect farmers' ability to substitute between groundwater and surface water.[1]

Armed with these estimates of groundwater demand elasticities, we quantify deadweight loss from the pumping cost externality. Deadweight loss only exists if farm i earns less economic surplus from pumping a marginal unit of groundwater than the total economic surplus lost by i's neighbors due to their marginally higher pumping costs. We parameterize constant elasticity demand curves for every farm in our sample, in order to calculate farm i's loss in surplus from consuming 1 acre-foot less. We then calculate the marginal increase in groundwater levels for all neighboring farms, resulting from farm i having extracted less.[2] Translating groundwater level increases into pumping cost decreases, we sum the marginal consumer surplus gained by each of farm i's neighbors—equivalent to farm i's short-run pumping cost externality. Our preliminary results suggest that the magnitude of this externality is relatively small (roughly \$2 per acre-foot on average) relative to these farms'

1. In ongoing work, we are incorporating data on farmers' surface water availability in order to estimate the elasticity of substitution between these two water sources.

2. In this draft, we simplify the hydrology and calculate effects in concentric circles around each extraction point. Future versions of this analysis will incorporate more realistic and sophisticated hydrogeological assumptions.

pumping costs (averaging $40 per acre-foot in our sample). We also find substantial variation both within and between California's groundwater basins. In fact, many farms in our sample likely contribute no deadweight loss, as their own marginal surplus loss is greater than the sum of their neighbors' marginal surplus gains. Our results suggest that a relatively small tax on groundwater extraction has the potential to eliminate millions of dollars of deadweight loss annually for California farmers.

This paper makes three main contributions. First, we provide the first large-scale estimates of the price elasticity of electricity demand in a major energy-using sector: agriculture in California, one of the most important agricultural sectors in the world. While many studies estimate the relationship between electricity prices and consumption in the residential sector (Alberini and Filippini (2011), Fell, Li, and Paul (2014), Ito (2014), Deryugina, MacKay, and Reif (2018)), far fewer studies focus on commercial and industrial electricity consumption (Paul, Myers, and Palmer (2009), Jessoe and Rapson (2015), Blonz (2016)). To the best of our knowledge, there exist no previous estimates of the price elasticity of electricity demand in the agricultural sector. By leveraging microdata from the universe of agricultural consumers in PGE's service territory, along with plausibly exogenous changes in farmers' marginal electricity price, we identify California farmers as relatively elastic electricity consumers.

Second, we estimate the price elasticity of groundwater demand for California farmers—a policy-relevant elasticity that has proven difficult to estimate due to both data and identification challenges. We overcome these challenges by combining comprehensive electricity consumption data with technical audits of groundwater pumps, and by leveraging exogenous variation in electricity prices (a major component of pumping costs) to credibly identify changes in farmers' effective price of groundwater. Many previous studies have estimated the price elasticity of water demand outside the agricultural sector (Hewitt and Hanemann (1995), Renwick and Green (2000), Olmstead, Hanemann, and Stavins (2007)), while others have focused specifically on groundwater demand in agriculture (Hendricks and Peterson (2012), Pfeiffer and Lin (2014), Badiani and Jessoe (2015), Mieno and Brozovic (2017)). Relative to this existing literature, we contribute well-identified groundwater demand estimates, for thousands of farms from one the most important agricultural regions in the world:

California's Central Valley. Our empirical strategy addresses a number of measurement and identification issues from the existing literature (noted in Mieno and Brozovic (2017)), by leveraging exogenous variation in pumping costs across time and space, in conjunction with novel data on electricity use, agricultural pump specifications, and groundwater levels.

Third, we use these groundwater elasticity estimates to compute the pumping cost externality for each farm in our sample, whereby farmers fail to internalize the effect of their own extraction on their neighbors' extraction costs. Because groundwater is a common-pool resource, it is associated with two important externalities: a "stock externality", as farmers do not fully internalize the full continuation value of the resource and extract too quickly, resulting in a "race to the bottom"; and a "pumping cost externality", in which farmer do not internalize how their own extraction lowers groundwater levels, thereby increasing the costs of extraction for neighboring farmers. We focus exclusively on the latter, an important externality that is often highlighted in theoretical work (Ostrom (1990), Provencher and Burt (1993)), but seldom estimated empirically. We leverage our quasi-experimental variation in electricity prices to estimate the "pumping cost" externality in California. Our preliminary results suggest that while the magnitude of the marginal external costs is likely small relative to the private costs of extracting groundwater, a small corrective tax could yield large welfare gains for Calfornia farmers.

This paper proceeds as follows. Section 2 provides background on groundwater pumping, California agriculture, and energy use in farming. Sections 3 and 4 describe our data and empirical strategy. In Section 5, we present results for our demand elasticity estimates, which we translate into deadweight loss calculations in Section 6. Section 7 concludes.

2 Background

2.1 Groundwater as a common-pool resource

Groundwater is a classic example of a common-pool (or open access) resource (Ostrom (1990)). Typical underground aquifers are large enough that excluding consumers is extremely costly. There are two main externalities that pertain to groundwater extraction, commonly known as the "stock externality" and the "pumping cost externality" (Provencher and Burt (1993)). The stock externality refers to the notion that the groundwater stock is finite, and as a result, when farmer i extracts from the aquifer in period t, this constrains other users in $t + 1$. While this is likely an important externality in practice, we do not address it in this paper because we lack the empirical machinery to do so.

The pumping cost externality, which is the focus of this paper, refers to the fact that when farmer i extracts an acre-foot of water from the aquifer, the water level in the aquifer drops. This, in turn, raises the cost of water extraction for other users of the aquifer. While farmer i will internalize the impact of his own extraction on his own costs of future extraction, he will not internalize the impacts of his extraction on the costs of his neighbors. This generates a classic externality, which is the focus of this paper. In particular, we aim to quantify the magnitude of the pumping cost externality – and the associated corrective Pigouvian tax – in the context of California agriculture.

2.2 Agriculture in California

California is a key player in global agricultural production. 17 percent of total U.S. crop value was produced in California in 2016 (Johnson and Cody (2015)). California farmers produce over 400 commodities, including close to half of all fruits, nuts, and vegetables grown in the United States, and there are may products for which California is the sole U.S. producer, such as almonds, artichokes, olives, and walnuts (California Department of Food and Agriculture (2011)). These goods are produced by over 77,000 farms in California.

Water is particularly important to California agriculture. Close to 80 percent of all of the state's water use each year takes place in the agricultural sector. Much of this water goes to crops. California has nearly 8.3 million harvested acres, 7.9 million of which were irrigated (Johnson and Cody (2015)). Farmers can either access surface water or groundwater. While these ultimately come from the same watershed, the rules governing use water differs.

Water use in California is governed by a complex system of rights.[3] Surface water rights fall into several categories, but all surface rights are heavily regulated; farmers using more water than allowed by their rights are subject to severe fines.

3. See Sawyers (2007) for more details.

By contrast, most agricultural groundwater rights are much more vague. The typical groundwater right in California is “overlying,” meaning that landowners whose property sits above an aquifer hav the right to extract the underlying groundwater.[4] The vast majority of groundwater use is unmetered, and users face no marginal costs of extraction beyond the energy costs of pumping (Bruno and Jessoe (2018)). This means that, so long as a farmer has property rights over land which sits above an aquifer, he can extract as much groundwater as he chooses. In the midst of the 2014 drought, in order to combat overextraction, California lawmakers passed the Sustainable Groundwater Management Act. This major piece of legislation, the first statewide regulatory scheme governing the extraction and use of groundwater, is aimed at enabling sustainable groundwater management via a series of local groundwater sustainability plans. However, these plans will not be finalized until 2020 at the earliest, leaving farmers free to extract at will.

2.3 Electricity for Pumping

Electricity is essential to modern groundwater extraction. The California Energy Commission reports that water use accounts for 19 percent of California’s electricity consumption, and close to 8 percent of the state’s energy is used on farms (California Energy Commission (2005)). The state’s investor-owned utilities spend nearly $50 million annually to improve energy efficiency in the agricultural sector. This makes water, and agricultural water use in particular, a major end-use of electricity in the state.

Our goal is to quantify the size of the pumping externality that farm i imposes on his neighbors. To do this, we exploit the fact that electricity is a major determinant of the cost of groundwater pumping to estimate farms’ price elasticity of demand for groundwater. Several previous papers have tried to estimate the price elasticity of groundwater demand using variation in energy costs (Hendricks and Peterson (2012), Pfeiffer and Lin (2014), Badiani and Jessoe (2015), and Mieno and Brozovic (2017)). However, Mieno and Brozovic (2017) point out that these papers may suffer from (non-classical) measurement error and identification

4. There are also “appropriative” groundwater rights, for users who do not own land above the aquifer, but these rights are lower-priority than the overlying rights. This means that users with appropriative water rights can only exercise these rights in the case of a surplus.

concerns which have the potential to bias the resulting estimates. Furthermore, the most detailed of these papers focus on small geographic locations where data were available; this increases the internal validity of these studies while reducing their external validity.

In contrast, we estimate the price elasticity of demand for agricultural groundwater throughout California's Central Valley, one of the most important agricultural regions in the world. We use a multi-step estimation process. We begin by using plausibly exogenous variation in energy pricing, both in the cross-section and over time to calculate the price elasticity of electricity consumption. We then combine this with detailed (time-varying) data on farm-specific pump characteristics and spatially explicit information on groundwater levels. With these additional data, we can estimate the price elasticity of groundwater demand with respect to (1) electricity prices, (2) groundwater levels, and finally, (3) water costs.

3 Data

3.1 Electricity Data

We begin by estimating how farmers' electricity consumption for groundwater pumping responds to changes in electricity price. To do this, we use confidential customer-level electricity datasets, which PGE's data management team prepared for us under a non-disclosure agreement. These data comprise the universe of agricultural electric consumers in PGE's service territory, and we observe each customer's monthly billing data at the service account level from 2008–2017. We aggregate service accounts up to 108,172 unique service points (i.e. the physical location of an electricity meter), allowing us to construct "monthified" panel of electricity consumption (in kWh) at the service point (SP) level.[5] We also observe several key covariates for each service point, including its latitude and longitude; an indicator for accounts with solar panels on net-energy metering (NEM), which we drop from our estimation sample; and an identifier to link service point locations to physical electricity meters.

5. PGE's monthly bill cycles are customer-specific, and billing periods typically do not end at the end of a calendar month. We "monthify" billed kWh for each service point by splitting/weight-averaging multiple bills in a single calendar month, in order to create a service point by month panel.

PGE's offers 23 distinct agricultural tariffs, and our billing data report the particular tariff associated with each monthly bill. Prices on each tariff update multiple times per year, and historic prices are publicly available, along with information on tariff-specific rules and eligibility criteria. This allows us to construct a 10-year panel of hourly volumetric (marginal) electricity prices, which we collapse to the monthly level by taking an unweighted average across hours. Then, we assign each service point a monthly average marginal electricity price in \$/kWh.

Table 1 describes each agricultural tariff in detail. Importantly, PGE classifies all agricultural consumers into 5 categories, based on their physical capital (i.e. pump size; internal combustion engine) and their type of electricity meter (i.e. conventional vs. smart meters).[6] While farmers may switch tariffs *within* categories, they may not switch tariffs *across* categories. This restriction lets us identify price elasticities of demand by instrumenting for farm i's marginal price with the marginal price of its within-category default tariff. Figure 1 plots time series of monthly average marginal prices for 3 of the most common tariffs in our estimation sample. Our identifying variation in electricity prices comes from (i) these time series not moving in parallel, and (ii) across-category tariff switching induced by PGE's smart meter rollout.

3.2 Pump Data

To complement our electricity data, we have rich data on agricultural groundwater pumps collected by PGE's Advanced Pumping Efficiency Program (APEP). These data include the universe of APEP-subsidized pump tests from 2011–2017, and we observe detailed measurements and technical specifications for 21,851 unique tests at 17,107 unique pump locations. Importantly, we also observe identifiers for the electricity meter associated with each pump test. This allows us to match pump tests to electricity service points, thereby isolating a sample of 11,851 service points for which agricultural groundwater pumping is confirmed to

6. Conventional meters record electricity consumption using an analog dial, whereas smart meters can digitally store the full time profile of consumption. During our sample period, PGE gradually phased out conventional meters, replacing them with smart meters capable of supporting time-varying electricity pricing.

be a major end-use.[7] We restrict our empirical analysis to this 11 percent subset of agricultural service points, in order to best isolate groundwater pumpers and avoid incorporating other agricultural electricity end-uses.[8]

Table 2 reports summary statistics for this subset of agricultural service points (in the right column). Compared to the full set of PGE's agricultural customers, APEP-matched service points tend to consume nearly twice as much electricity, and tend to pay lower marginal prices. Only 24 percent of service points shift across tariff categories, and the vast majority of switches are triggered by PGE's smart meter rollout. Figure 2 reveals that APEP-matched service points are heavily concentrated in California's Central Valley, and appear to be a geographically representative subset of PGE's agricultural customers.

Besides helping us identify a subset of agricultural consumers who pump groundwater, APEP data allow us to characterize pump-specific groundwater production functions. Groundwater extracted is Leontief in electricity (for pumps with electric motors), and 1 kWh of electricity in will yield a particular volume of groundwater out (measured in acre-feet (AF)). This kWh per AF relationship is governed by physics:

$$\frac{\text{kWh}}{\text{AF}} = \text{kW} \div \frac{\text{AF}}{\text{hour}} = \frac{[\text{Lift (feet)}] \times [\text{Flow (gallon/minute)}]}{[\text{Operating pump efficiency (\%)}] \times [\text{Constant}]} \div \frac{\text{AF}}{\text{hour}} \tag{1}$$

The power (kW) required to pump 1 acre-foot is directly proportionate to both the vertical distance the water must travel to the surface (i.e. lift) and the speed at which the water travels (i.e. flow). It is inversely proportionate to rate at which the pump converts electric energy into the movement of water (i.e. operating pump efficiency). We can simplify (1) by converting from gallons to acre-feet:

$$\Rightarrow \quad \frac{\text{kWh}}{\text{AF}} = \frac{[\text{Lift (feet)}] \times [\text{Constant}]}{\text{Operating pump efficiency (\%)}} \tag{2}$$

7. Pumping is likely the *only* end-use at most matched service points, as PGE typically installs a dedicated meters for each pump.

8. We are currently working on using satellite images to predict whether service points outside the APEP-matched sample are also groundwater pumpers. We hope to incorporate these farms into future analysis, as there are likely many groundwater pumps that never received an APEP-subsidized pump test.

For each APEP pump test, we observe measurements of kWh/AF, operating pump efficiency, flow, and lift. We also observe the standing water level, or the baseline depth of water in the absence of pumping. Because pumping lowers the water level at a given location, standing water levels help us more accurately calibrate how changes in aquifer depth impact lift for each pump.[9]

3.3 California Water Data

While a given farm's pumping technology is relatively constant in the short run, its kWh/AF conversion factor is sensitive to short-run changes in standing water levels. In order to capture these short-run shocks in pumping costs, we use publicly available groundwater data from California's Department of Water Resources collected under the California Statewide Groundwater Elevation Monitoring (CASGEM) Program.[10] These data report depth below the surface at 16,015 unique monitoring stations during our 2008–2017 sample period, with an average of 25 measurements at each location at different points in time. We rasterize all measurements within each month (and quarter), using inverse distance weighting to interpolate a two-dimensional surface of average depth at each point in space. This allows us to construct a monthly (and quarterly) panel of estimated groundwater depths at each service point in our APEP-matched sample.

We assign each service point to both a groundwater basin and a water district, using publicly available shapefiles maintained by the California Department of Water Resources.[11] Groundwater basins are broadly defined by stratigraphic barriers through which water does not travel horizontally. We control for annual changes in water levels that impact all farms within the same water basin. Water districts are administrative entities that govern farmers'

9. Lift is (approximately) the sum of the standing water level, drawdown (i.e., how much pump i impacts its own depth), and other pump-specific factors (e.g., discharge pressure, gauge corrections, height of the pump above the surface). Drawdown depends on rate of extraction (i.e. flow) and the physical properties of the substrata. Greater flow increases drawdown, as water levels fall with faster extraction. More transmissive (or porous) rock formations have lower drawdown, because water levels are able to reequilibrate horizontally more quickly.

10. These data are available at: `https://water.ca.gov/Programs/Groundwater-Management/Groundwater-Elevation-Monitoring--CASGEM`

11. Water basin shapefiles are available at `https://water.ca.gov/Programs/Groundwater-Management/Bulletin-118`. Water district shapefiles are available at `https://data.cnra.ca.gov/dataset/water-districts`.

annual allocations of surface water. Because groundwater and surface water are obvious substitutes, we control for annual shocks to farms' surface water allocations at the water district level. This helps to isolate variation in pumping behavior driven by variation in pumping costs, rather than variation in the availability of groundwater substitutes.[12]

3.4 Groundwater Prices and Quantities

We combine the above data sources into a panel of groundwater prices and quantities, at the service point by month level. To convert from electricity (kWh or \$/kWh) to groundwater (AF or \$/AF), we simply need to populate a kWh/AF conversion factor for every panel observation. We construct estimates for kWh/AF by parameterizing Equation (2) using (i) monthly (or quarterly) rasters of groundwater depths at each service point; (ii) pump-specific conversions between standing water level and lift, as calculated from APEP pump tests; and (iii) APEP-measured operating pump efficiencies. We take unweighted averages of APEP variables across multiple pumps within a single service point; we also extrapolate each service point's first pump test backwards, extrapolate its last pump test forwards, and interpolate between multiple pump tests using a triangular kernel in time.

Table 3 reports summary statistics for this merged panel dataset. We observe 3.45 unique pump tests for the average APEP-matched service point, and APEP data reveal substantial (cross-sectional) variation in operating pump efficiencies and kWh/AF conversion factors. Our constructed kWh/AF estimates tend to moderate extremely values, which compresses the right tail kWh/AF values (while also slightly shifting this distribution left). Interestingly, implied marginal groundwater prices exhibit far less seasonal variation than marginal electricity prices. This is because groundwater levels tend to be relatively higher in summer months (compared to winter months), which tends to reduce (estimated) kWh/AF in months when electricity prices are highest.

12. We are currently working to incorporate additional spatial data products, including (i) shapefiles of common land units, which correspond to agricultural fields; (ii) shapefiles of parcels, which roughly correspond to farms; and (iii) the USDA Cropland Data Layer, which uses satellite imagery to classify fields by crop type. Together, these data will enable us to aggregate up from service points to farms, and to build a panel of crop choice at the farm-year level.

4 Empirical Strategy

This section outlines our empirical strategy for estimating farmers' demand for groundwater pumping. First, we estimate price elasticities of demand for electricity, for the full sample of agricultural consumers where we can match an electricity meter to a groundwater pump. Next, we estimate price elasticities of demand for groundwater, by translating prices/quantities of electricity into prices/quantitites of water using data on (i) technical pumping production functions and (ii) groundwater depths across space and time.

4.1 Electricity Demand

We estimate monthly electricity demand using the following specification:

$$\sinh^{-1}\left(Q_{it}^{\text{elec}}\right) = \beta \log\left(P_{it}^{\text{elec}}\right) + \gamma_i + \delta_t + \varepsilon_{it} \tag{3}$$

The dependent variable is kWh consumed at electricity service point i in month t, transformed using the inverse hyperbolic sine function (which closely approximates the natural log transformation but includes zero in its support).[13] P_{it}^{elec} is unit i's marginal electricity price (in \$/kWh), averaged across all hours in month t. We include unit by month-of-year fixed effects (γ_i), in order to control for within-pump/month average consumption (e.g., service point i in March). We also include month-of-sample fixed effects (δ_t) to control for trends in both electricity prices (which rise over time) and pumping behavior. Alternative specifications include groundwater basin by year fixed effects (to control for annual between-basin variation in groundwater depth), water district by year fixed effects (to control for annual shocks to surface water allocations), and unit-specific linear time trends. We two-way cluster standard errors by service point and month-of-sample, which accommodates both arbitrary within-unit serial correlations and arbitrary spatial correlations with each monthly cross-section.

In order to identify the demand elasticity in Equation (3), we must purge any endogenous variation in unit i's marginal electricity price. PGE's agricultural tariff *schedules* are

13. Since 15 percent of observations in this panel are zeroes, we apply the inverse hyperbolic sine transformation to avoid dropping months where farms consume zero kWh for groundwater pumping.

the outcome by statewide regulatory proceedings, and marginal prices are linear in kWh consumed.[14] While an individual farmer cannot plausibly influence how PGE sets prices, most farmers may select between alternative tariff schedules—effectively choosing which marginal electricity price they face. PGE restricts this choice to be within 5 tariff categories defined by farmers' physical capital (e.g. pump size and type) and type of electric meter (i.e. conventional vs. "smart" meters). We instrument for unit i's marginal price with the marginal price on the default tariff *within* unit i's category (i.e. bolded tariffs in Table 1), which eliminates selection bias from a high-volume pumper choosing a tariff with advantageously low volumetric prices.

Farmers may also shift *across* tariff categories, which could similarly bias our elasticity estimates. If such a shift reflects a change in physical pumping capital—for example, upgrading from a < 35 hp pump to a ≥ 35 hp pump—then a change in marginal price (or within-category marginal default price) would coincide with a mechanical increase in electricity consumption. We control for such endogenous changes in price by interacting unit fixed effects with a categorical variable for the 3 types physical capital that define PGE tariff categories: small pumps, large pumps, and internal combustion engines. On the other hand, if unit i shifted across categories because PGE replaced its conventional meter with a smart meter, we would not expect such a shift to coincide with any other changes to unit i's pumping behavior.[15] Hence, meter-induced shifts in tariff categories are unlikely to lead to *endogenous* changes in unit i marginal electricity price. We also instrument with *lagged* default prices to purge potential endogeneity in the timing of unit i's smart meter installation.

14. By contrast, PGE's residential electricity tariffs use increasing block pricing, where a household's marginal price is endogenous to its own consumption (Ito (2014)). Linear marginal prices simplify our estimation of agricultural electricity demand, because farm i's marginal price is determined *solely* by its tariff schedule.

15. During our 2008–2017 sample period, PGE gradually installed smart meters for the vast majority of its customers. The timing of PGE's smart meter rollout was determined by institutional and geographic factors, which were outside of customers' control. Previous research has established that PGE did not specify the rollout to target customers with specific usage patterns (Blonz (2016)).

4.2 Groundwater Demand

We seek to estimate the causal effect of groundwater price on groundwater consumption, and this demand elasticity is linearly approximated by the coefficient β:

$$\log\left(Q_{it}^{\text{water}}\right) = \beta \log\left(P_{it}^{\text{water}}\right) \tag{4}$$

We construct Q_{it}^{water} and P_{it}^{water} using the *estimated* conversion factor $\widehat{\frac{\text{kWh}}{\text{AF}}}_{it}$, which has measurement error and is also potentially endogenous. Hence, the same measurement error and endogeneity is present on both the left-hand side and the right-hand side of Equation (4). We can rewrite this expression decomposing $\widehat{\frac{\text{kWh}}{\text{AF}}}_{it}$ on both sides:

$$\log\left(Q_{it}^{\text{elec}}\right) - \log\left(\widehat{\frac{\text{kWh}}{\text{AF}}}_{it}\right) = \beta\left[\log\left(P_{it}^{\text{elec}}\right) + \log\left(\widehat{\frac{\text{kWh}}{\text{AF}}}_{it}\right)\right] \tag{5}$$

Rearranging:

$$\log\left(Q_{it}^{\text{elec}}\right) = \beta \log\left(P_{it}^{\text{elec}}\right) + (\beta + 1)\log\left(\widehat{\frac{\text{kWh}}{\text{AF}}}_{it}\right) \tag{6}$$

This expression is algebraically equivalent to Equation (4), but it isolates the endogenous estimated conversion factor in one right-hand-side variable. We estimate an analogous regression specification:

$$\sinh^{-1}\left(Q_{it}^{\text{elec}}\right) = \beta^{\text{e}} \log\left(P_{it}^{\text{elec}}\right) + (\beta^{\text{w}} + 1)\log\left(\widehat{\frac{\text{kWh}}{\text{AF}}}_{it}\right) + \gamma_i + \delta_t + \varepsilon_{it} \tag{7}$$

This specification is similar to Equation (3), except that we can now interpret β^{e} and β^{w} as the price elasticity of demand for groundwater. We allow this elasticity to vary depending on the source of variation in pumping costs—groundwater depths may be more salient to farmers than electricity prices, or vice versa.[16] As in the electricity regressions, we purge electricity price endogeneity by instrumenting P_{it}^{elec} with within-category default prices (see description above).

To identify β^{w}, we must overcome three potential sources of bias. First, farmers may choose to alter their pumping technologies in order to change $\widehat{\frac{\text{kWh}}{\text{AF}}}_{it}$, and such changes are

16. A strict Neoclassical interpretation would assume $\beta^{\text{e}} = \beta^{\text{w}}$, as the optimizing farmer should respond to all short-run changes in P_{it}^{water} identically.

likely correlated with Q_{it}^{elec}. Second, $\widehat{\frac{\text{kWh}}{\text{AF}}}_{it}$ is a function of unit i's groundwater depth, which is mechanically linked to Q_{it}^{elec}—when unit i consumes electricity to extract groundwater, its localized groundwater level falls, thereby increasing $\widehat{\frac{\text{kWh}}{\text{AF}}}_{it}$. Third, $\widehat{\frac{\text{kWh}}{\text{AF}}}_{it}$ incorporates measurement error both from interpolating rasterized groundwater depths across space and from interpolating/extrapolating unit i's APEP measurements across time.

We instrument for $\log\left(\widehat{\frac{\text{kWh}}{\text{AF}}}_{it}\right)$ using logged groundwater depth averaged across unit i's full groundwater basin.[17] This purges potential endogeneity driven by changes in pumping technologies, and eliminates bias induced by measurement error in unit i's pump specifications in month t. It also breaks the mechanical relationship between $\widehat{\frac{\text{kWh}}{\text{AF}}}_{it}$ and Q_{it}^{elec}, as farm i's extraction should have a negligible contemporaneous effect on average groundwater levels across the whole basin. Finally, instrumenting with basin-wide average depth mitigates measurement error from having spatially interpolated groundwater measurements into a (potentially overfit) gridded raster.

While Equation (7) isolates the shared endogenous component of Q^{water} and P^{water} on the right-hand side, a more standard approach would be to estimate:

$$\sinh^{-1}\left(Q_{it}^{\text{water}}\right) = \beta \log\left(P_{it}^{\text{water}}\right) + \gamma_i + \delta_t + \varepsilon_{it} \tag{8}$$

We also estimate Equation (8), instrumenting for $\log\left(P_{it}^{\text{water}}\right)$ using the logged average marginal electricity price of unit i's within-category default tariff. This instrument isolates changes in the effective price of groundwater driven *only* by plausibly exogenous changes in the marginal electricity price. It also removes measurement in $\widehat{\frac{\text{kWh}}{\text{AF}}}_{it}$ from the right-hand side, which prevents measurement error on the left-hand side from biasing our point estimates.[18]

17. We instrument with groundwater depth in logs (rather than levels) because logging both sides of Equation (2) implies that $\log(\text{kWh}/\text{AF}_{it})$ is linear in $\log(\text{lift})$, and a percentage change in depth should yield a similar percentage change in lift.

18. In most cases, classical measurement error on the left-hand side does not bias point estimates. However, consider the regression $(Y_i + \eta_i) = \beta(X_i + \omega_i) + \varepsilon$, where η_i and ω_i each denote classical measurement error, and where $\text{Cov}(\eta_i, \omega_i) \neq 0$. After conditioning on $(X_i + \omega_i)$, the remaining measurement error on the left-hand side is no longer (conditionally) classical, and could bias $\hat{\beta}$. In Equation (8), measurement error from $\widehat{\text{kWh}/\text{AF}}_{it}$ enters directly on the right-hand side and inversely on the left-hand side. Instrumenting with default electricity prices removes this correlation between left-hand- vs. right-hand-side measurement error.

5 Results

5.1 Electricity Demand

Table 4 reports results from estimating Equation 3. Column (1) does not instrument for marginal electricity price, yielding an unidentified point estimate for $\hat{\beta}$. Column (2) instruments using unit i's within-category default marginal price, which eliminates bias from farmers choosing their own electricity tariffs. The direction of this bias is not obvious *ex ante*, because farmers are choosing between tariffs with both volumetric (\$/kWh) and fixed (\$/kW) price components.[19] Comparing Columns (1) vs. (2), we see that on average, farmers with higher electricity consumption tend to choose tariffs with relatively lower fixed charges per kW and relatively higher prices per kWh.

Column (3) eliminates the other source of price endogeneity, by interacting unit fixed effects with indicators for (i) small pumps ($<$ 35 hp), (ii) large pumps ($\geq$ 35 hp), and (iii) auxiliary internal combustion engines. While only 5 percent of units shift across tariff categories due to changes in their physical capital, the resulting simultaneous changes in Q_{it}^{elec} and P_{it}^{elec} induce substantial bias in Column (2) point estimate. Column (3) reports our preferred estimate of -1.17, after having purged both sources of price endogeneity.

The magnitude of this elasticity estimate is surprisingly large, considering that electricity demand tends to be extremely inelastic in other contexts. The literature on electricity demand has focused heavily on the residential sector, and recent estimates have found elasticities of -0.08 to -0.48 in the short run (Reiss and White (2005); Alberini and Filippini (2011); Fell, Li, and Paul (2014)) and -0.09 to -0.73 in the medium to long run (Alberini and Filippini (2011); Ito (2014); Deryugina, MacKay, and Reif (2018)).[20] Fewer estimates exist for commercial or industrial electricity demand. Paul, Myers, and Palmer (2009) es-

19. PGE tariffs with relatively high volumetric (i.e. marginal) prices tend to have relatively low fixed prices, and vice versa. Two farmers with the same average electricity consumption may optimally choose different tariffs. Suppose farmer A operates a 300 hp pump for 50 hours per month, while farmer B operates 50 hp pump for 300 hours per month. Farmer A should prefer a low fixed price and a high volumetric price, while farmer B should prefer a high fixed price and a low marginal price.

20. These estimates uses monthly or annual variation in electricity prices, which aligns with our empirical strategy. Other studies leverage hourly variation in electricity prices have estimated electricity demand elasticities ranging from -0.03 to -0.25 (Wolak (2011); Jessoe and Rapson (2014); Fowlie et al. (2018); Ito, Ida, and Tanaka (2018)).

timate commercial/industrial elasticities of -0.11 to -0.16 in the short run, and -0.29 to -0.40 in the long run. Jessoe and Rapson (2015) find no demand response to dynamic pricing in these sectors, while Blonz (2016) estimates elasticities of -0.08 to -0.22 using hourly price variation for PGE's small commercial/industrial customers. To our knowledge, we provide the first large-scale estimates of electricity demand elasticities in the agricultural sector.

Columns (4)–(6) report three additional elasticity estimates, each intended to assuage any remaining concerns over price endogeneity. Column (4) includes separate year fixed effects for each water basin and each water district, to control for potential time-varying confounders related to water depth or surface water availability. The resulting point estimate of -1.02 is quite similar, albeit slightly attenuated. Column (5) instruments with the 6- and 12-month lags of the default price (rather than the contemporaneous default price), to account for potential endogeneity in the timing of PGE's smart meter rollout.[21] This yields a nearly identical point estimate, implying that farmers' electricity consumption did not meaningfully change in anticipation of a smart meter installation. Finally, Column (6) adds 11,175 unit-specific linear time trends, to confirm that we are not identifying $\hat{\beta}$ solely off of monotonic trends in price and quantity. The resulting point estimate of -0.76 is attenuated, as linear trends remove much of the (good) variation in electricity prices over time. Even so, we still find a tightly estimated elasticity that is substantially larger than virtually all previous estimates for the elasticity of electricity demand.

5.2 Groundwater Demand

Table 5 presents our results for estimating farmers' groundwater demand. Each column estimates Equation (7) using our preferred strategy for identifying the elasticity with respect to electricity price: instrumenting for $\log\left(P_{it}^{\text{elec}}\right)$ with within-category default prices, and interacting unit fixed effects with indicators for each category of physical pumping capital. Note that we report $\hat{\beta}^{\text{e}}$ and $\hat{\beta}^{\text{w}}$, where the latter subtracts 1 from the regression coefficient on $\log\left(\widehat{\frac{\text{kWh}}{\text{AF}}}_{it}\right)$. We interpret each coefficient as the elasticity of demand for groundwater

21. Recall that farmers may shift across tariff categories (inducing changes to their within-category default price) due to *either* changes in their physical capital *or* the installation of a smart meter.

with respect to one component of the price of groundwater, holding the other component constant.

Column (2) reports our preferred estimates of $\hat{\beta}^{\mathrm{e}}$ and $\hat{\beta}^{\mathrm{w}}$, where we instrument for $\log\left(\widehat{\frac{\mathrm{kWh}}{\mathrm{AF}}}_{it}\right)$ with logged groundwater depth in month t averaged across unit i's full groundwater basin. Comparing $\hat{\beta}^{\mathrm{w}}$ in Columns (2) vs. (1), instrumenting with average depth appears to alleviate bias due to measurement error in $\log\left(\widehat{\frac{\mathrm{kWh}}{\mathrm{AF}}}_{it}\right)$.[22] The exclusion restriction requires that unit i's pumping behavior have no contemporaneous impact on basin-wide average groundwater depths. Such feedback effects between the dependent variable and the instrument would be extremely unlikely for three reasons: (i) unit i is small relative to the geographic footprint of its groundwater basin; (ii) thousands of other pumpers are also extracting from the same basin; (iii) basin-wide average groundwater levels do not instantaneously reequilibrate after extraction at one point in space. Column (3) restricts the sample to the 3 largest groundwater basins, each of which has over 1,000 units in our estimation sample.[23] The resulting $\hat{\beta}^{\mathrm{w}}$ estimate is quite similar, which should assuage concerns that the instrument is invalid due to a few large farms located in very small groundwater basins.

The magnitudes of our $\hat{\beta}^{\mathrm{e}}$ estimates are quite similar to results from the electricity-only regressions, especially comparing $\hat{\beta}^{\mathrm{e}} = -1.21$ from Column (1) of Table 5 the analogous estimate of $\hat{\beta} = -1.17$ from Column (3) of Table 4. Perhaps surprisingly, $\hat{\beta}^{\mathrm{e}}$ is also nearly identical to our instrumented $\hat{\beta}^{\mathrm{w}}$ estimates. This implies that a 1 percent change in the effective price of groundwater has the *same* effect on farmers' pumping behavior, whether that change comes via their marginal electricity price or via their pump's kWh/AF conversion factor. It also suggests that farmers are quite attentive to their true costs of pumping, and that they reoptimize their pumping behavior identically in response to either type of price variation—as Neoclassical theory would predict.

Similar to our elasticity estimates for electricity, our groundwater elasticity estimates are also quite large relative to most previous estimates. Recent studies have also exploited

22. We discuss three potential sources of bias in β^{w} in Section 4.2 — (i) endogenous changes to pumping technologies, (ii) the mechanical relationship between extraction and depth at a given location, and (iii) measurement error. Bias from (i) or (ii) appears unlikely, as they should bias our β^{w} *away from* zero, rather than towards zero.

23. These basins are the San Joaquin Valley, the Sacramento Valley, and the Salinas Valley. The number of agricultural groundwater pumpers in each basin is likely much larger, as our estimation sample comprises only the subset of PGE customers that we can confident match to an APEP-subsidized pump test.

variation in energy prices, but yielding far smaller magnitudes: Hendricks and Peterson (2012) find an elasticity of -0.10, and Pfeiffer and Lin (2014) find an elasticity of -0.27 (both for agricultural groundwater in Kansas). Bruno and Jessoe (2018) estimate demand elasticities of -0.17 to -0.22 within the Coachella Valley of California, which is a unique setting where groundwater extraction is directly priced. Previous studies have also estimated farmers' elasticity of demand for surface water, most notably Hagerty (2018), who finds an elasticity of -0.23 for surface water in California agriculture. While estimates of surface water demand often range as high as -0.8 in specific locations (Schoengold, Sunding, and Moreno (2014); Hagerty (2018)), we find agricultural groundwater demand to be much more elastic.[24] Substitution between groundwater and surface water is likely a major factor explaining the large magnitudes of our elasticities estimates.

Columns (4)–(6) report three alternate versions of our preferred estimates in Column (2). First, to account for the inherent tradeoff between spatial density vs. temporal frequency of groundwater measurements, Column (4) re-estimates Equation (7) using groundwater data rasterized at the quarterly (rather than monthly) level. Whereas our preferred monthly rasters are able to capture groundwater measurements at greater temporal frequency, quarterly rasters have greater accuracy in the cross-section by incorporating more distinct measurement sites. The resulting $\hat{\beta}^{\mathrm{w}}$ estimate increases in magnitude, however the average depth instrument has less predictive power at the (coarser) quarterly level. Column (5) includes water basin by year and water district by year fixed effects, yielding only slightly attenuated point estimates despite eliminating much of the variation in the average depth instrument. In Column (6), we instrument with 6- and 12-month lags of average depth (rather than contemporaneous depth), as it is possible (albeit unlikely) that farmers pump less in months with lower groundwater levels for some reason other than pumping costs. These lagged instruments substantially increase $\hat{\beta}^{\mathrm{e}}$ and $\hat{\beta}^{\mathrm{w}}$; however, the small first stage F-statistic indicates a weak instrument, and we interpret these results with caution.

Table 6 reports results from estimating Equation (8), with groundwater quantity as the dependent variable, and instrumenting for the composite groundwater price with default

24. Estimates for urban water demand have found similar elasticities, ranging from -0.10 to -0.76 (Nataraj (2011); Ito (2013); Baerenklau, Schwabe, and Dinar (2014); Wichman (2014);Buck et al. (2016); Wichman, Taylor, and Haefen (2016); Hagerty (2018)).

electricity prices. While these estimates identify $\hat{\beta}$ using *only* variation in default electricity prices, we demonstrate above that farmers respond almost identically to variation in *either* component of their effective groundwater price. The resulting point estimates are quite similar to, but slightly smaller than, our electricity demand estimates. Interestingly, combining P_{it}^{water} into a single regressor removes much of the variation used to estimate separate coefficients in Table 5. This is because electricity prices and groundwater depths are seasonally correlated—groundwater levels are lowest (making pumping more expensive) in the winter months, when electricity prices are also low. This likely explains why $\hat{\beta}$ estimates in Table 6 are smaller than $\hat{\beta}^{\text{e}}$ and $\hat{\beta}^{\text{w}}$ estimates in Table 5.

6 Open-Access Externality

We use our elasticity estimates to calculate the "pumping cost" externality for each farm in our sample. This calculates the deadweight loss created when farm i fails to internalize how its own groundwater extraction increases its neighbors' costs of extraction. Importantly, farm i only creates deadweight loss if its own surplus gained from pumping a marginal acre-foot is *less* than the sum of surplus lost by *all* neighboring farms due to marginally lower groundwater levels. We focus exclusively on the (more static) "pumping cost" externality, where marginal changes in farm i's extraction induce marginal changes neighboring farms' groundwater levels and pumping costs. We do not investigate the (more dynamic) "stock externality," where farmers' failure to internalize the full continuation value of groundwater stocks causes them to deplete the stock too quickly.

Figure 3 illustrates our conceptual framework for the simplified case where four farms share a single aquifer. Each farm has a unique groundwater demand curve, which governs the amount of groundwater it extracts (Q_i) at a given marginal pumping cost (P_i). If Farm 1 extracts less groundwater, it incurs a loss in consumer surplus (i.e., the red triangle). However, this has the effect of increasing groundwater levels for Farms 2, 3, and 4—which leads to farm-specific decreases in their marginal pumping costs (since the translation from depth to \$/AF depends on each farm's unique electricity price and pumping technology).[25] At lower pumping costs, Farms 2, 3, and 4 increase extraction and capture additional consumer surplus (i.e., the green trapezoids). If the sum of surplus gained by Farms 2, 3, and 4 is greater than the surplus lost by Farm 1, then Farm 1's open-access extraction is creating deadweight loss.

We require two inputs to operationalize Figure 3 and calculate the open-access externality for each farm i. First, we must parameterize i-specific groundwater demand curves. To do this, we assume constant elasticity of demand, with a homogeneous elasticity of $\epsilon = -1.12$ (based on our preferred estimate from Table 6). Each demand curve is simply $Q = A_i P^{\epsilon}$, where A_i is calibrated to unit i's groundwater price and quantity. For this exercise, we focus on June–July of 2016, the two months with the greatest groundwater extraction in our last complete year of data.

Second, we must characterize how a marginal decrease in farm i's groundwater extraction affects groundwater levels at each neighboring farm j. The requires assumptions for both the spatial patterns of how groundwater moves horizontally through the substrata (i.e. at what distance from farm i do groundwater levels respond to i's extraction) and the speed at which groundwater levels reequilibrate (i.e. how quickly does farm i's extraction impact farm j's groundwater levels). We make two crude assumptions for the sake of tractability: (i) when farm i extracts 1 acre-foot less, this causes groundwater levels to increase equally at all points within a radius of r miles, where the increase is equal to 1 foot divided by the area (in acres) of the r-mile circle;[26] (ii) when farm i extracts 1 acre-foot less in June, groundwater levels throughout the r-mile circle fully reequilibrate by July. We acknowledge that these assumptions are not ideal, and future versions of this analysis will incorporate more realistic assumptions based on detailed hydrogeological models of California's groundwater basins (Sunding, Zilberman, and Brozovic (2010)).

Table 7 reports deadweight loss calculations for each farm (i.e. service point) in the three most common groundwater basins in our sample: the Sacramento Valley (641 APEP-

25. The increase in groundwater levels may also vary across farms sharing an aquifer, as hydrogeological factors govern the speed that groundwater flows both horizontally and vertically through the substrata.

26. Aquifers are not bathtubs, and groundwater levels can vary substantially within a given r-mile circle. For this simple exercise, we assume that groundwater levels increase in parallel throughout the circle. This preserves baseline heterogeneity in groundwater levels across space.

matched service points actively pumping in June 2016), the Salinas Valley (941), and the San Joaquín Valley (5,325). In each basin, the median farm would have lost about \$0.50 in consumer surplus from pumping 1 AF less in June 2016 (i.e. the red triangle from Figure 3). This consumer surplus loss varies substantially, as the 1st-quartile farm would have lost up to \$1.96, while the 3rd-quartile farm would have lost as little as \$0.16.[27] The bottom three panels calculate the total change in consumer surplus, including both farm i's lost surplus and the surplus gained by all neighboring farms j within an r-mile radius. We present two calculations for each groundwater basin and radius: (i) "APEP" columns calculate $\sum_j \Delta CS_j$ for all neighboring service points in our APEP-matched estimation sample; (ii) "Scaled" columns inflate the number of neighbors by the proportion of *unmatched* PGE service points on agricultural tariffs.[28] Bolded numbers indicate positive welfare changes from reducing deadweight loss—where farm i's marginal open-access externality is greater than its own private marginal surplus.

Four notable patterns emerge from the distributions of changes in deadweight loss. First, greater spatial density of groundwater pumping increases the magnitude of the open-access externality, as $\sum_j \Delta CS_j$ sums across a greater number of neighbors. This is not surprising, yet Table 7 highlights stark differences across basins and between "APEP" vs. "Scaled" calculations. The static pumping externality is largest in the San Joaquín Valley, which contains the majority of all PGE agricultural customers. Accounting for unmatched groundwater pumpers is also important, as the APEP matched sample dramatically under-counts the number of neighbors impacted by farm i's extraction.

Second, increasing the size of the radius r has an ambiguous effect on the size of the externality: larger circles include more neighbors but spread the additional 1 AF of groundwater across a greater area, yielding smaller cost reductions for each neighbor. Table 7 revels that larger cirlces tend to yield more positive welfare changes, especially in the San Joaquín

27. This variation in ΔCS_i reflects baseline differences in P_i and Q_i, which we use to parameterize each i-specific demand curve.

28. Our econometric analysis only includes PGE service points that we can directly match to an APEP pump test. However, this sample represents a relatively small subset of PGE agricultural customers—most of whom are also groundwater pumpers (groundwater pumping is the leading agricultural energy end-use in California). PGE's agricultural tariffs are typically more expensive than its non-agricultural tariffs, hence farmers have little incentive to pay agricultural tariffs for non-agricultural end uses.

Valley. Where pumping is the most geographically dense, the number of farm i's neighbors increases fast enough with r to offset increasingly marginal changes in groundwater depth.

Third, we observe a substantial amount of variation in deadweight loss within each groundwater basin, radius, and sample. While the 3rd quartile of deadweight loss is positive for all nine "Scaled" calculations, the 1st quartile is almost always negative. This implies that over 25 percent of farms in each groundwater basin are likely contributing deadweight loss on the margin, but a different 25 percent of farms are likely *not* contributing any marginal deadweight loss. An optimal price- or quantity-based policy would need to account for variation in the open-access externality imposed by each farm i, though there is less cross-sectional variation in the marginal externality itself ($\sum_j \Delta CS_j$) than in the marginal deadweight loss ($\Delta CS_i + \sum_j \Delta CS_j$).

Fourth, the magnitudes of our deadweight loss calculations are small relative to farmers' marginal pumping costs. For the 5,325 farmers in the San Joaquín Valley, their average marginal cost of groundwater pumping in June 2016 was \$60 per AF. Our largest estimates of the marginal deadweight loss are an order of magnitude smaller. At the same time, these 5,325 farmers extracted nearly 400,000 AF of groundwater in June 2016 alone. This suggests that due to farmers' relatively elastic demand for groundwater, a proportionately small tax to internalize the (static) open-access externality could have eliminated nearly \$1 million in deadweight loss, for just one summer month. Also, Table 7 likely understates the true magnitudes of deadweight loss by only considering ΔCS_j for July 2016. In reality, if farm i pumps 1 AF less in June, neighboring farms should continue to accrue marginal increases in consumer surplus after July.

These calculations are preliminary, and we interpret them with caution. Once we incorporate more sophisticated hydrogeological assumptions, we plan to (i) replace simple radii with more spatially nuanced representations of each farm's neighbors; (ii) relax our assumption of a parallel increase in depth for all neighboring farms, to allow for groundwater equilibration to occur at different speeds for different neighbors; and (iii) relax our crude treatment of the timing, to allow neighbors to accrue losses in surplus for as many months as the short-run change in depth is likely to persist. We also plan to apply machine learning techniques to predict which unmatched farms are groundwater pumpers, rather than simply scaling by the rate of unmatched service points. As an input to this prediction exercise, we will assign all service points to polygons for Common Land Units (i.e. agricultural fields) and parcels (i.e. contiguous farms under common ownership), to remove the subset of unmatched services points located far from agricultural crop lands.

Assigning services points to polygons will also let us aggregate up to the more natural cross-sectional unit—the farm. The will prevent us from mistaking multiple pumps on the same farm for neighbors. We also plan to overlay the USDA's Cropland Data Layer, in order to assign a specific crop (or set of crops) to each farm. This will allow us to estimate heterogeneous demand elasticities by crop type. While our current calculations are driven entirely by differences in A_i and variation in the spatial density of i's neighbors, incorporating heterogeneous ϵ_i will allow us to capture differences between annuals vs. perennials, crops with high vs. low water intensity, and high- vs. low-value crops. For example, almond farmers (perennial, high water intensity, high value) are likely much more inelastic groundwater users than alfalfa farmers (perennial, high water intensity, low value), due to almonds' higher marginal value product per acre-foot. It follows that alfalfa farmers are more likely to impose large pumping externalities on almond farmers than vice versa.

7 Conclusion

This paper estimates the price elasticities of demand both for agricultural electricity consumption (for groundwater pumping) and for agricultural groundwater consumption. We overcome common data challenges by linking customer-specific electricity consumption data with data from in-person pump efficiency audits, in order construct accurate measures for farmers' groundwater consumption and pumping costs. We credibly identify our elasticity estimates using exogenous variation in agricultural electricity tariffs, and our results reveal farmers to be much more elastic groundwater pumpers than previously thought. Using these elasticity estimates, we quantify the "pumping cost" externality—or the extent to which farm i increases its neighbors' groundwater extraction costs by removing a marginal unit of water from their shared aquifer. Our preliminary results suggest that the magnitude of this externality is likely small relative to the private costs of groundwater pumping, and that a small corrective tax could yield substantial welfare gains for California farmers.

These results are still preliminary, and future versions of this paper will incorporate more sophisticated hydrogeological assumptions to more accurately characterize the spatial nature of the pumping cost externality. We also plan to use detailed geospatial data in order to aggregate up from electricity service points to farms, as well as high-resolution satellite imagery to classify farms by crop type. This will allow us to estimate heterogeneous elasticities by crop type, and to incorporate these heterogeneous elasticities into our externality calculations. A richer version of Table 7 will differentiate between annuals vs. perennials, crops with high vs. low water intensity, and high- vs. low-value crops. This will allow us to evaluate the extent to which the pumping cost externality contributes to misallocation in crop choice. We also plan to use quasi-experimental variation in farmers' pumping capital to estimate the long-run elasticity of demand for groundwater.

References

Alberini, Anna, and Massimo Filippini. 2011. "Response of residential electricity demand to price: The effect of measurement error." *Energy Economics* 33:889–895.

Badiani, Reena, and Katrina Jessoe. 2015. "The impact of electricity subsidies on groundwater extraction and agricultural production." Working paper.

Baerenklau, Kenneth A., Kurt A. Schwabe, and Ariel Dinar. 2014. "The residential water demand effect of increasing block rate water budgets." *Land Economics* 90 (4): 683–699.

Blonz, Joshua A. 2016. "Making the best of the second-best: Welfare consequences of time-varying electricity prices." Energy Institute at Haas Working Paper 275.

Bruno, Ellen, and Katrina Jessoe. 2018. "Water markets and climate change adaptation: Micro-level evidence on agricultural water demand." Mimeo.

Buck, Steven, Maximilian Auffhammer, Stephen Hamilton, and David Sunding. 2016. "Measuring welfare losses from urban water supply disruptions." *Journal of the Association of Environmental and Resource Economists* 3 (3): 743–778.

California Department of Food and Agriculture. 2011. "California agricultural resource directory 2010–2011."

California Energy Commission. 2005. "California's water-energy relationship." CEC Report 700-2005-011-SF.

Deryugina, Tatyana, Alexander MacKay, and Julian Reif. 2018. "The long-run dynamics of electricity demand: Evidence from municipal aggregation." Working paper.

Fell, Harrison, Shanjun Li, and Anthony Paul. 2014. "A new look at residential electricity demand using household expenditure data." *International Journal of Industrial Organization* 33:37–47.

Fowlie, Meredith, Catherine Wolfram, C. Anna Spurlock, Annika Todd, Patrick Baylis, and Peter Cappers. 2018. "Default effects and follow-on behavior: Evidence from an electricity pricing program." Working paper.

Hagerty, Nick. 2018. "Liquid constrained in California: Estimating the potential gains from water markets." Working paper.

Hendricks, Nathan P., and Jeffrey M. Peterson. 2012. "Fixed effects estimation of the intensive and extensive margins of irrigation water." *Journal of Agricultural and Resource Economics* 37 (1): 1–19.

Hewitt, J.A., and W.M. Hanemann. 1995. "A discrete/continuous choice approach to residential water demand under block rate pricing." *Land Economics* 71 (2): 173–192.

Ito, Koichiro. 2013. "How do consumers respond to nonlinear pricing? Evidence from household water demand." Working paper.

———. 2014. "Do consumers respond to marginal or average price? Evidence from nonlinear electricity pricing." *American Economic Review* 104 (2): 537–563.

Ito, Koichiro, Takanori Ida, and Makoto Tanaka. 2018. "Moral suasion and economic incentives: Field experimental evidence from electricity demand." *American Economic Journal: Economic Policy* 10 (1): 240–267.

Jessoe, Katrina, and David Rapson. 2014. "Knowledge is (less) power: Experimental evidence from residential electricity use." *American Economic Review* 104 (4): 1417–1438.

———. 2015. "Commercial and industrial demand response under mandatory time-of-use electricity pricing." *Journal of Industrial Economics* 63 (3): 397–421.

Johnson, Renee, and Betsy A. Cody. 2015. "California agricultural production and irrigated water use." Congressional Research Service report.

Mieno, Taro, and Nicholas Brozovic. 2017. "Price Elasticity of Groundwater Demand: Attenuation and Amplification Bias Due to Incomplete Information." *American Journal of Agricultural Economics* 99 (2): 401–426.

Nataraj, Shanthi. 2011. "Does marginal price matter? A regression discontinuity approach to estimating water demand." *Journal of Environmental Economics and Management* 61:198–212.

Olmstead, S.M., Michael W. Hanemann, and R.N. Stavins. 2007. "Water demand under alternative price structures." *Journal of Environmental Economics and Management* 54 (2): 181–198.

Ostrom, Elinor. 1990. *Governing the Commons: The Evolution of Institutions for Collective Action.* Cambridge University Press.

Paul, Anthony, Erica Myers, and Karen Palmer. 2009. “A partial adjustment model of U.S. electricity demand by region, season, and sector.” Resources for the Future Working Paper DP 08-50.

Pfeiffer, Lisa, and C.-Y. Cynthia Lin. 2014. “The effects of energy prices on agricultural groundwater extraction from the High Plains aquifer.” *American Journal of Agricultural Economics* 96 (5): 1349–1362.

Provencher, Bill, and Oscar Burt. 1993. “The Externalities Associated with the Common Property Exploitation of Groundwater.” *Journal of Environmental Economics and Management* 24:139–158.

Reiss, Peter C., and Matthew W. White. 2005. “Household electricity demand, revisited.” *Review of Economic Studies* 72:853–883.

Renwick, E.M., and R.D. Green. 2000. “Do residential water demand side management policies measure up? An analysis of eight California water agencies.” *Journal of Environmental Economics and Management* 40 (1): 37–55.

Sawyers, Gary. 2007. “A primer on California water rights.” UC Davis Cooperative Extension working paper.

Schoengold, Katrina, David L. Sunding, and Georgina Moreno. 2014. “Price elasticity reconsidered: Panel estimation of an agricultural water demand function.” *Water Resources Research* 42.

Sunding, David L., David Zilberman, and Nicholas Brozovic. 2010. “On the spatial nature of the groundwater pumping externality.” *Resource and Energy Economics* 32 (2): 154–164.

Wichman, Casey J. 2014. “Perceived price in residential water demand: Evidence from a natural experiment.” *Journal of Economic Behavior and Organization* 107:308–323.

Wichman, Casey J., Laura O. Taylor, and Roger H. von Haefen. 2016. “Conservation policies: Who responds to price and who responds to prescription?” *Journal of Environmental Economics and Management* 79:114–134.

Wolak, Frank A. 2011. “Do residential customers respond to hourly prices? Evidence from a dynamic pricing experiment.” *American Economic Review* 101 (3): 83–87.

Tables and Figures

Figure 1: Average Marginal Electricity Prices

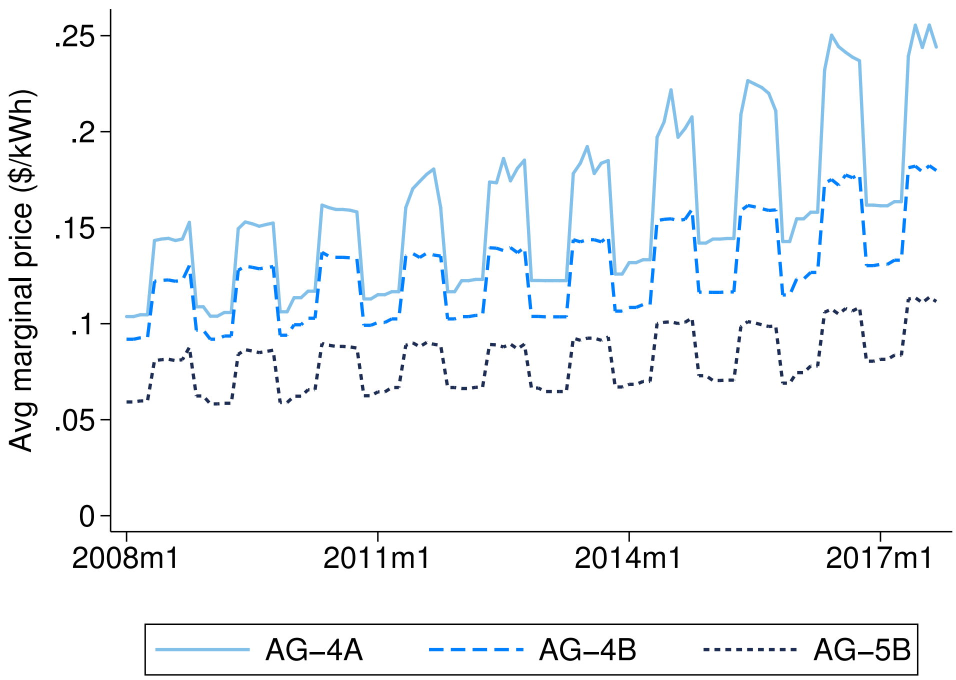

Notes: This figure reports the times series of monthly average marginal electricity prices ($/kWh) for three of the most common agricultural tariffs. Prices are systematically higher during summer months (May–October). Much of our identifying variation in monthly electricity prices comes from these monthly price times series not rising perfectly in parallel.

Figure 2: PGE Agricultural Customers

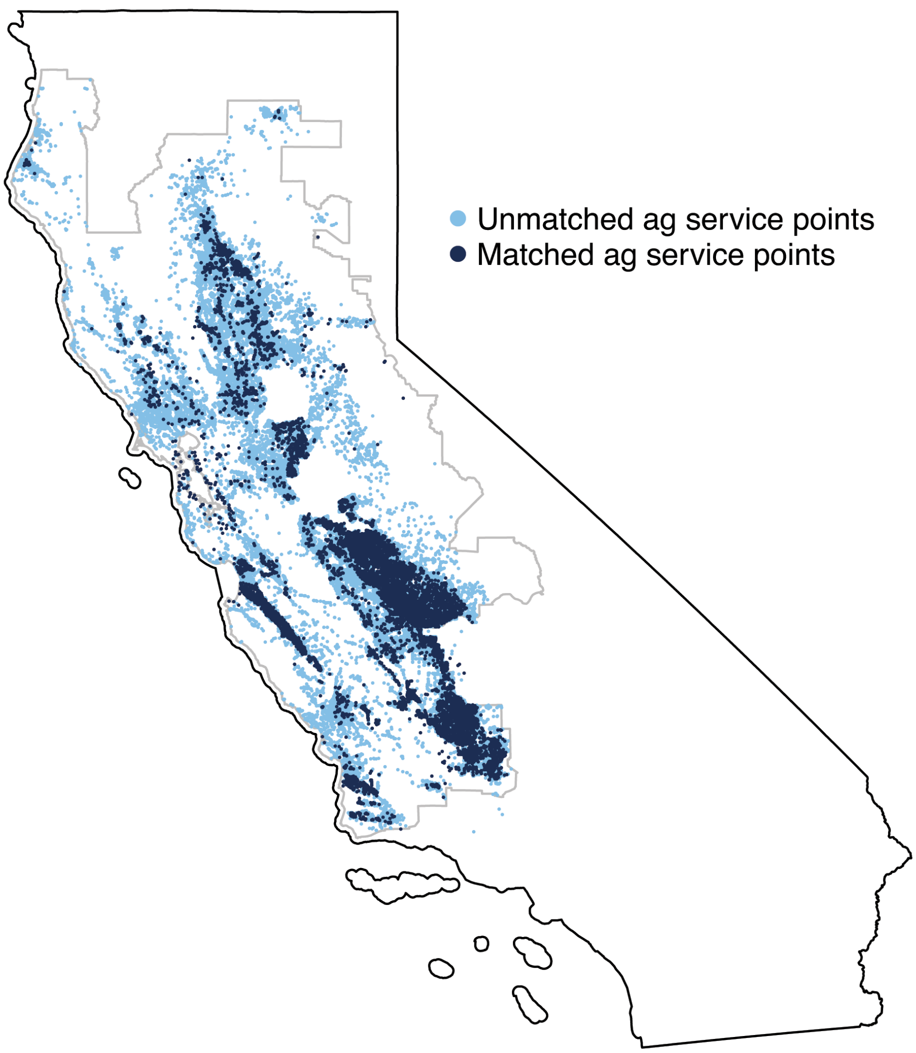

Notes: This figure maps the locations of all agricultural service points served by PGE. Dark blue dots indicate the 11,851 service point that we can match directly to an APEP pump test. Light blue dots indicate unmatched agricultural service points. The light grey outline is the geographic boundary of PGE's service territory.

Figure 3: Framework for Calculating the Open-Access Externality

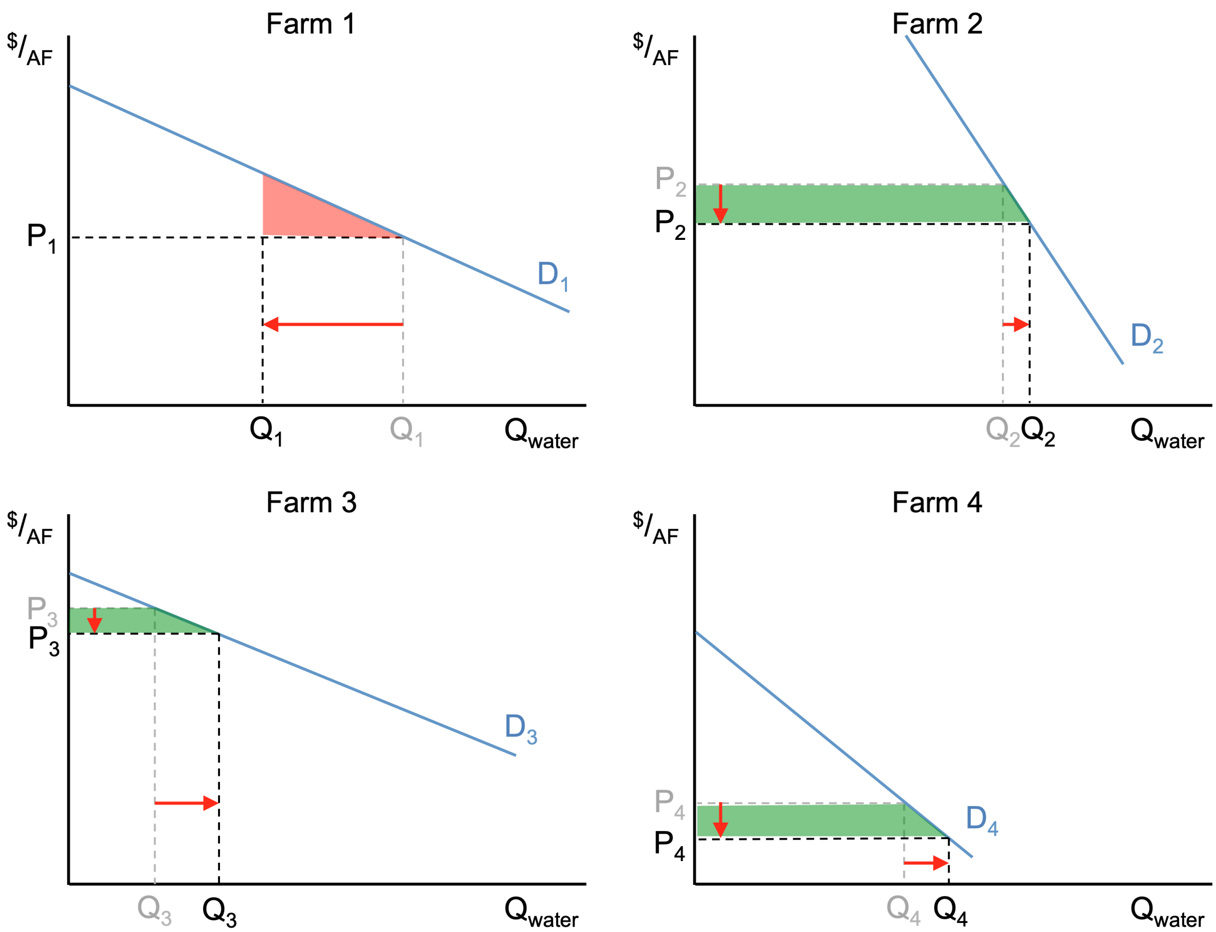

Notes: This figure illustrates how we conceptualize and calculate the open access pumping externality that farm i imposes on its neighbors—by not internalizing how its own groundwater extraction increases extraction costs for other nearby farms. In example, Farm 1 has three neighbors (Farms 2, 3, and 4), and all four farms have distinct demand curves for groundwater. If Farm 1 pumps less than its desired quantity of groundwater, its consumer surplus falls (illustrated by the red triangle). However, when Farm 1 extracts less groundwater, this slightly increases groundwater levels for neighboring Farms 2, 3, and 4. Because Farms 2, 3, and 4 have distinct pumping technologies and electricity prices, a parallel increase in groundwater levels leads to idiosyncratic decreases in their effective prices of groundwater. Farms 2, 3, and 4 respond to these price decreases by slightly increasing groundwater extraction, leading to increased consumer surplus (illustrated by the green trapezoids). If the sum of the (green) surplus gains for all of Farm 1's neighbors is greater than Farm 1's own (red) surplus loss, then the Social Planner can increase welfare by taxing Farm 1's groundwater extraction.

Table 1: PGE Agricultural Tariffs

Category	Tariff	Description	Percent
Small pumps, conventional meters single motor <35 hp, or multiple motors summing to <15 hp	**1A**	High price per kWh (not time-varying), fixed charge per hp connected	3.0
Large pumps, conventional meters single motor ≥ 35 hp, or multiple motors summing to ≥ 15 hp, or single overloaded motor ≥ 15 hp	**1B**	High price per kWh (not time-varying), fixed charge per max kW consumed	8.1
Small pumps, smart meters single motor <35 hp, or multiple motors summing to <15 hp	**4A** (4D)	High prices per kWh (higher in peak hours), fixed charges per hp connected, very high peak prices on 14 summer Event Days	7.2
	5A (5D)	Lower prices per kWh (peak & offpeak), no Event Day price increases, higher fixed charges per hp	2.7
	RA (RD)	Lower peak prices per kWh, higher off-peak prices per kWh, no Event Day price increases, choice between MTW or WTF peak days	1.2
	VA (VD)	Lower peak prices per kWh, higher off-peak prices per kWh, no Event Day price increases, choice of 3 shorter 4-hour peak periods	0.9
Large pumps, smart meters single motor ≥ 35 hp, or multiple motors summing to ≥ 15 hp, or single overloaded motor ≥ 15 hp	**4B** (4E)	High prices per kWh (higher in peak hours), fixed charges per max kW consumed	20.1
	5B (5E)	Much lower prices per kWh (peak & offpeak), higher fixed charge per max kW	37.8
	4C (4F)	Slightly lower prices per kWh (peak & offpeak), higher fixed charges per kW shifted to peak, very high peak prices on 14 summer Event Days	2.4
	5C (5F)	Much lower prices per kWh (peak & offpeak), higher fixed charges per kW shifted to peak, very high peak prices on 14 summer Event Days	7.8
	RB (RE)	Higher prices per kWh (peak & off-peak), choice between MTW or WTF peak days, lower fixed charges per max kW (in summer)	1.5
	VB (VF)	Higher prices per kWh (peak & off-peak), choice of 3 shorter 4-hour peak periods, lower fixed charges per max kW (in summer)	0.6
Customers transitioning off internal combustion engines	**ICE**	Very low price per kWh (high in peak hours), fixed charge per max kW consumed	6.8

Notes: This table provides a rough summary of PGE's 23 electricity tariffs for agricultural customers. The first column lists the 5 disjoint categories of customers, defined (primarily) by physical pumping capital and electricity meters. Effective default tarrifs within each group are in bold, and farmers may switch tariffs *within* a category (but not *across* categories). All tariffs have fixed (\$/kW) and volumetric (\$/kWh) prices that vary by summer vs. winter. All time-of-use tariffs (i.e. all but 1A and 1B) also vary between peak (12:00pm–6:00pm on summer weekdays), partial peak (8:30am–9:30pm on weekends), and off-peak periods. DEF tariffs are functionally equivalent to their ABC analogs, and are holdovers for the earliest customers to adopt time-of-use pricing. Actual tariffs are far more complex, and tariff documents are available at `https://www.pge.com/tariffs/index.page`. The right-most column reports the percent of observations in our main estimation sample on each tariff.

Table 2: Summary Statistics – Electricity Data

	All Ag Customers	Matched to Pumps
Service point-month observations	9,991,458	1,168,553
Unique service points (SPs)	108,172	11,851
SPs that switch tariff categories	44,414	2,844
SPs that switch categories (pumping capital)	3,454	561
SPs that switch categories (smart meters)	43,045	2,553
Share of SP-months on time-varying tariffs	0.702	0.886
Share of SP-months on peak-day tariffs	0.295	0.152
Monthly electricity consumption (kWh)	6080.9	12055.7
	(39783.1)	(25075.1)
Monthly electricity consumption (kWh), summer	8249.6	17589.1
	(45660.8)	(29818.5)
Monthly electricity consumption (kWh), winter	3849.7	6362.8
	(32498.8)	(17232.5)
Average marginal electricity price ($/kWh)	0.148	0.113
	(0.050)	(0.042)
Average marginal electricity price ($/kWh), summer	0.171	0.130
	(0.051)	(0.044)
Average marginal electricity price ($/kWh), winter	0.126	0.096
	(0.037)	(0.032)
Average monthly bill ($, non-zero bills)	936.66	1814.15
	(4662.71)	(3285.26)
Average monthly bill ($, non-zero bills), summer	1398.90	2821.16
	(5847.34)	(4020.99)
Average monthly bill ($, non-zero bills), winter	456.17	764.99
	(2888.86)	(1742.21)

Notes: The left column reports summary statistics for the universe of agricultural electricity customers in PGE service territory, from 2008–2017. The right column includes the subset of agricultural customers that we successfully match to a groundwater pump in the APEP pump test dataset—i.e., our main estimation sample. "Pumping capital" denotes tariff category switches driven by shifts between small pumps ($<$ 35 hp) and large pumps ($\geq$ 35 hp), or adding/removing an auxiliary internal combustion engine. Most tariff category switches were driven by PGE's smart meter rollout. Time-varying tariffs (i.e. all except 1A and 1B) have higher marginal prices during peak demand hours. Peak-day tariffs (i.e. 4A, 4D, 4C, 4F, 5C, 5F) have very high marginal prices during peak hours on the 14 highest-demand summer days. Monthly bills include both volumetric ($/kWh) and fixed charges ($/kW, $/hp, and $/day). Summer months are May–October. Standard deviations of sample means in parentheses.

Table 3: Summary Statistics – Pump Tests and Groundwater Consumption

	Matched to Pumps
Service point-month observations	1,168,553
Unique service points (SPs)	11,851
Matched APEP points per SP	3.45
	(8.80)
Operating pump efficiency (%)	54.46
	(11.52)
kWh per AF conversion factor (APEP measured)	430.30
	(254.21)
kWh per AF conversion factor (constructed)	346.94
	(206.03)
Monthly groundwater consumption (AF)	49.0
	(151.7)
Monthly groundwater consumption (AF), summer	74.0
	(191.9)
Monthly groundwater consumption (AF), winter	23.4
	(86.6)
Average marginal groundwater price ($/AF)	39.91
	(29.63)
Average marginal groundwater price ($/AF), summer	41.76
	(31.34)
Average marginal groundwater price ($/AF), winter	38.01
	(27.63)

Notes: These summary stats are from the merged panel of groundwater prices and quantities, which combines electricity data, pump test data, and groundwater data. We observe 3.45 unique APEP pump tests for the average matched service point, although 37 percent of service points match to only a single APEP test. Our constructed kWh per AF conversion factor (i.e. $\widehat{\text{kWh/AF}}_{it}$) uses monthly groundwater rasters to capture changes in (measured) kWh per AF over time, and estimation error compresses the right tail of distribution of measured kWh per AF. Monthly groundwater consumption divides electricity consumption (kWh) by $\widehat{\text{kWh/AF}}_{it}$. Grounwater prices multiply marignal electricity prices ($/kWh) by $\widehat{\text{kWh/AF}}_{it}$. Summer months are May–October. Standard deviations of sample means in parentheses.

Table 4: Estimated Demand Elasticities – Electricity

	(1) OLS	(2) IV	(3) IV	(4) IV	(5) IV	(6) IV
$\log\left(P_{it}^{\text{elec}}\right)$	-1.31^{***}	-1.58^{***}	-1.17^{***}	-1.02^{***}	-1.18^{***}	-0.76^{***}
	(0.11)	(0.17)	(0.16)	(0.14)	(0.21)	(0.17)
Instrument(s):						
Default $\log\left(P_{it}^{\text{elec}}\right)$		Yes	Yes	Yes		Yes
Default $\log\left(P_{it}^{\text{elec}}\right)$, lagged					Yes	
Fixed effects:						
Unit × month-of-year	Yes	Yes	Yes	Yes	Yes	Yes
Month-of-sample	Yes	Yes	Yes	Yes	Yes	Yes
Unit × physical capital			Yes	Yes	Yes	Yes
Water basin × year				Yes		
Water district × year				Yes		
Unit-specific linear time trends						Yes
Service point units	11,175	11,175	11,175	11,121	10,924	11,175
Months	117	117	117	117	105	117
Observations	1.05M	1.05M	1.05M	1.04M	0.91M	1.05M
First stage F-statistic		4136	7382	7508	757	4776

Notes: Each regression estimates Equation (3) at the service point by month level, where the dependent variable is the inverse hyperbolic sine transformation of electricity consumed by service point i in month t. We estimate IV specifications via two-stage least squares, instrumenting with either unit i's within-category default logged electricity price in month t *or* the 6- and 12-month lags of this variable. "Physical capital" is a categorical variable for (i) small pumps, (ii) large pumps, and (iii) internal combustion engines, and unit × physical capital fixed effects control for shifts in tariff category triggered by the installation of new pumping equipment. Water basin × year fixed effects control for broad geographic trends in groundwater depth. Water district × year fixed effects control for annual variation in surface water allocations. All regressions drop solar NEM customers, customers with bad geocodes, and months with irregular electricity bills (e.g. first/last bills, bills longer/shorter than 1 month, overlapping bills for a single account). Standard errors (in parentheses) are clustered by service point and by month-of-sample. Significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table 5: Estimated Demand Elasticities – Groundwater

	(1) IV	(2) IV	(3) IV	(4) IV	(5) IV	(6) IV
$\log\left(P^{\text{elec}}_{it}\right)$: $\hat{\beta}^{\text{e}}$	-1.21^{***}	-1.39^{***}	-1.38^{***}	-1.26^{***}	-1.23^{***}	-1.68^{***}
	(0.17)	(0.18)	(0.19)	(0.17)	(0.16)	(0.21)
$\log\left(\widehat{\frac{\text{kWh}}{\text{AF}}}_{it}\right)$: $\hat{\beta}^{\text{w}}$	-0.92^{***}	-1.37^{***}	-1.32^{***}	-1.72^{***}	-1.24^{***}	-2.04^{***}
	(0.11)	(0.25)	(0.28)	(0.30)	(0.24)	(0.45)
Instrument(s):						
Default $\log\left(P^{\text{elec}}_{it}\right)$	Yes	Yes	Yes	Yes	Yes	Yes
log (Avg depth in basin)		Yes	Yes	Yes	Yes	
log (Avg depth in basin), lagged						Yes
Fixed effects:						
Unit × month-of-year	Yes	Yes	Yes	Yes	Yes	Yes
Month-of-sample	Yes	Yes	Yes	Yes	Yes	Yes
Unit × physical capital	Yes	Yes	Yes	Yes	Yes	Yes
Water basin × year					Yes	
Water district × year					Yes	
Groundwater time step	Month	Month	Month	Quarter	Month	Month
Only basins with > 1000 SPs			Yes			
Service point units	10,159	10,121	9,324	10,134	10,086	9,890
Months	117	116	116	117	116	105
Observations	0.93M	0.83M	0.80M	0.89M	0.83M	0.70M
First stage F-statistic	6932	129	144	61	69	20

Notes: Each regression estimates Equation (7) at the service point by month level, where the dependent variable is the inverse hyperbolic sine transformation of electricity consumed by service point i in month t. We report estimates for $\hat{\beta}^{\text{e}}$ and $\hat{\beta}^{\text{w}}$, where the latter subtracts 1 from the estimated coefficient on $\log\left(\widehat{\text{kWh}/\text{AF}}_{it}\right)$. We estimate IV specifications via two-stage least squares, and all regressions instrument for P^{elec}_{it} with unit i's within-category default logged electricity price in month t (consistent with our preferred specification from Table 4). We instrument for $\log\left(\widehat{\text{kWh}/\text{AF}}_{it}\right)$ with either logged average groundwater depth across unit i's basin, or the 6- and 12-month lags of this variable. "Physical capital" is a categorical variable for (i) small pumps, (ii) large pumps, and (iii) internal combustion engines, and unit × physical capital fixed effects control for shifts in tariff category triggered by the installation of new pumping equipment. Water basin × year fixed effects control for broad geographic trends in groundwater depth. Water district × year fixed effects control for annual variation in surface water allocations. Column (3) restricts the sample to only the three most common water basins (San Joaquin Valley, Sacramento Valley, and Salinas Valley), each of which contains over 1000 unique SPs in our estimation sample. Column (4) uses a quarterly panel of groundwater depths to construct $\log\left(\widehat{\text{kWh}/\text{AF}}_{it}\right)$ and the instrument, rather than a monthly panel. All regressions drop solar NEM customers, customers with bad geocodes, months with irregular electricity bills (e.g. first/last bills, bills longer/shorter than 1 month, overlapping bills for a single account), and pumps with implausible test measurements. Standard errors (in parentheses) are clustered by service point and by month-of-sample. Significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table 6: Estimated Demand Elasticities – Groundwater

	(1) OLS	(2) IV	(3) IV	(4) IV	(5) IV	(6) IV
$\log(P_{it}^{\text{water}})$	-0.81^{***}	-1.12^{***}	-1.16^{***}	-1.12^{***}	-0.97^{***}	-1.14^{***}
	(0.09)	(0.15)	(0.17)	(0.15)	(0.14)	(0.21)
Instrument(s):						
Default $\log(P_{it}^{\text{elec}})$		Yes	Yes	Yes	Yes	
Default $\log(P_{it}^{\text{elec}})$, lagged						Yes
Fixed effects:						
Unit × month-of-year	Yes	Yes	Yes	Yes	Yes	Yes
Month-of-sample	Yes	Yes	Yes	Yes	Yes	Yes
Unit × physical capital	Yes	Yes	Yes	Yes	Yes	Yes
Water basin × year					Yes	
Water district × year					Yes	
Groundwater time step	Month	Month	Month	Quarter	Month	Month
Only basins with > 1000 SPs			Yes			
Service point units	10,159	10,159	9,342	10,159	10,118	9,926
Months	117	117	117	117	117	105
Observations	0.93M	0.93M	0.85M	0.93M	0.93M	0.82M
First stage F-statistic		2835	2735	2846	4633	486

Notes: Each regression estimates Equation (8) at the service point by month level, where the dependent variable is the inverse hyperbolic sine transformation of groudnwater consumed by service point i in month t. We estimate IV specifications via two-stage least squares, and Columns (2)–(5) instrument for P_{it}^{water} with unit i's within-category default logged electricity price. Column (6) instruments with the 6- and 12- month lags of this variable. "Physical capital" is a categorical variable for (i) small pumps, (ii) large pumps, and (iii) internal combustion engines, and unit × physical capital fixed effects control for shifts in tariff category triggered by the installation of new pumping equipment. Water basin × year fixed effects control for broad geographic trends in groundwater depth. Water district × year fixed effects control for annual variation in surface water allocations. Column (3) restricts the sample to only the three most common water basins (San Joaquin Valley, Sacramento Valley, and Salinas Valley), each of which contains over 1000 unique SPs in our estimation sample. Column (4) uses a quarterly panel of groundwater depths to construct both Q_{it}^{water} and P_{it}^{water}, rather than a monthly panel. All regressions drop solar NEM customers, customers with bad geocodes, months with irregular electricity bills (e.g. first/last bills, bills longer/shorter than 1 month, overlapping bills for a single account), and pumps with implausible test measurements. Standard errors (in parentheses) are clustered by service point and by month-of-sample. Significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table 7: Deadweight Loss Calculations for June 2016

	Sacramento Valley		Salinas Valley		San Joaquín Valley	
APEP units (i)	641		964		5,325	
ΔCS_i from 1 AF less						
25th percentile	−1.05		−1.61		−1.96	
50th percentile	−0.45		−0.59		−0.56	
75th percentile	−0.16		−0.28		−0.22	
10-mile radius	APEP	Scaled	APEP	Scaled	APEP	Scaled
Mean # of neighbors (j)	112	2,235	257	950	363	2,581
$\Delta CS_i + \sum_j \Delta CS_j$						
25th percentile	−1.03	−0.67	−1.54	−1.35	−1.86	−1.27
50th percentile	−0.43	−0.08	−0.50	−0.30	−0.45	**0.10**
75th percentile	−0.14	**0.15**	−0.18	**0.08**	−0.10	**0.66**
20-mile radius	APEP	Scaled	APEP	Scaled	APEP	Scaled
Mean # of neighbors (j)	236	4,729	498	1,842	1,047	7,430
$\Delta CS_i + \sum_j \Delta CS_j$						
25th percentile	−1.02	−0.36	−1.51	−1.18	−1.67	−0.31
50th percentile	−0.42	**0.14**	−0.45	−0.13	−0.28	**1.25**
75th percentile	−0.13	**0.46**	−0.12	**0.31**	**0.07**	**2.09**
30-mile radius	APEP	Scaled	APEP	Scaled	APEP	Scaled
Mean # of neighbors (j)	339	6,775	710	2,625	1,736	12,322
$\Delta CS_i + \sum_j \Delta CS_j$						
25th percentile	−1.01	−0.07	−1.45	−1.02	−1.47	**0.42**
50th percentile	−0.41	**0.39**	−0.41	**0.04**	−0.11	**2.59**
75th percentile	−0.12	**0.74**	−0.07	**0.51**	**0.27**	**3.68**

Notes: This table reports calculations for the open-access externality, for service point i in our main estimation sample located in the three largest groundwater basins (Sacramento Valley, Salinas Valley, and San Joaquín Valley). This simple exercise includes only service points with positive groundwater extraction in June 2016, with counts reported in the top row. For each unit i, we calculate their private decrease in consumer surplus from pumping 1 AF less in June, 2016 (reported in the top panel). Next, we translate unit i's 1-AF lower extraction in June into the corresponding marginal increase in groundwater levels in July across i-centered circles of radius r. For neighboring units j within each circle, this marginal increase in groundwater levels translates to a marginal decrease in j's effective price of groundwater. The bottom three panels report the net effect on *total* consumer surplus, subtracting i's lost consumer surplus in June from increased consumer surplus in July summed across all units j. Bolded numbers indicate positive welfare changes, consistent with unit i imposing a negative open-access externality greater than its own private benefit. "APEP" columns include only neighbors in our APEP-matched estimation sample, which almost certainly understates the magnitude of $\sum_j \Delta CS_j$ (by summing over only a subset of nearby agricultural groundwater pumpers). "Scaled" columns inflate the number of i's neighbors based on the ratio of match-to-unmatched PGE agricultural service points in each groundwater basin. We calculate all changes in consumer surplus by parameterizing unit-specific groundwater demand curves, imposing a homogeneous and constant elasticity of $\epsilon = -1.12$ (based on our estimate from Table 6, Column (2)). All units are in \$/AF.